Repeated Delegation

by *Elliot Lipnowski & João Ramos*[*]

NEW YORK UNIVERSITY

May 2015

Abstract

A principal sequentially delegates project adoption decisions to an agent, who can assess project quality but has lower standards than the principal. In equilibrium, the principal allows bad projects in the future to incentivize the agent to be selective today. The optimal contract, termed Dynamic Capital Budgeting, comprises two regimes. First, the principal provides an expense account to fund projects and yields full discretion to the agent. The account accrues interest until hitting a cap. While the account grows, the agent is willingly selective. After enough projects, the second regime begins, and the agent loses his autonomy forever.

JEL codes: C73, D23, D73, D82, D86, G31

[*]Email: elliot.lipnowski@nyu.edu and ramos@nyu.edu

This work has benefited from discussions with Heski Bar-Isaac, V. Bhaskar, Adam Brandenburger, Sylvain Chassang, Brendan Daley, Joyee Deb, Ross Doppelt, Ignacio Esponda, Eduardo Faingold, Johannes Hörner, R. Vijay Krishna, Vijay Krishna, Laurent Mathevet, David Pearce, Debraj Ray, Alejandro Rivera, Ariel Rubinstein, Evan Sadler, Maher Said, Tomasz Sadzik, Ennio Stacchetti, Alex Wolitzky, Chris Woolnough, Sevgi Yuksel, and seminar participants at New York University, EconCon, University of Waterloo, Midwest Economic Theory, and Canadian Economic Theory Conference. The usual disclaimer applies.

1 Introduction

Many economic activities are arranged via delegated decision making. In practice, those with the necessary information to make a decision may differ—and, indeed, have different interests—from those with the legal authority to act. Such relationships are often ongoing, consisting of many distinct decisions to be made over time, with the conflict of interest persisting throughout. A state government that funds local infrastructure may be more selective than the local government equipped to evaluate its potential benefits. A university bears the cost of a hired professor, relying on the department to determine candidates' quality. The Department of Defense funds specialized equipment for each of its units, but must rely on those on the ground to assess their need for it. Our focus is on how such repeated delegation should optimally be organized, and on how the relationship evolves over time.

Beyond the absence of monetary incentives, formal contingent contracting may be difficult for two reasons. First, it may be impractical for the informed party to produce verifiable evidence supporting its recommendations. Second, it might be unrealistic for the controlling party to credibly cede authority in the long run. Even so, the prospect of a future relationship may align the actors' interests: both parties may be flexible concerning their immediate goals, with a view to a healthy relationship.

Even for decisions that are, in principle, separate—in which today's course of action has no bearing on tomorrow's prospects—the controlling party can connect them, as a means to discipline the informed party now. If the university restricts the physics department to ten hires per decade, this might persuade the physics department to be discerning in the present, as hiring a mediocre physicist would crowd out a good one. By employing a *budgeting* rule, the controlling party imposes a cost on the agent for excessive spending, better aligning their interests.

We study an infinitely repeated game between a principal ("she") with full authority over a decision to be made in an uncertain world; she relies on an agent ("he") to assess the state. Each period, the principal must choose whether or not to initiate a project, which may be good (i.e. high enough value to offset its cost) or bad. The principal herself is ignorant of the current project's quality, but the agent knows it. The players have partially aligned preferences: both prefer a good project to any other outcome, but they disagree on which projects are worth taking. The principal wishes to fund only good projects, while the agent always prefers to invest in any project. For instance, consider the ongoing relationship between local and state governments. Each year, a county can request state government funds for the construction of a park. The state, taking into account past funding decisions, decides whether or not to fund it. The park would surely benefit the county, but the state must weigh this benefit against the money's opportunity cost. To assess this trade-off, the state relies on the county's local expertise. We focus on the case in which the principal needs the agent: the ex-ante expected value of a project is not enough to offset its cost. If the county were never selective in its proposals, the state would never want to fund the park. The agent's private information is transient: project types are independent across time, and a given project affects only within-period payoffs.[1]

To delegate—to cede control at the ex-ante stage—entails some vulnerability. If our principal wants to make use of the agent's expertise, she must give him the liberty to act. In funding a park, the state government risks wasting taxpayer money. Acting on a county's recommendation, the state won't know whether the park is truly valuable to the community, even after it is built. Furthermore, if the state makes a policy of funding each and every park the county requests, then it risks wasting a lot of money on many unneeded parks. This vulnerability limits the freedom that the agent can expect from the principal in the future. The state government cannot credibly reward a county's fiscal restraint today by promising carte blanche in the future.

The present conflict of interest would be resolved if the principal could sell permanent control to the agent.[2] In keeping with our leading applications, we focus on the repeated interaction without monetary transactions.[3] The Department of

[1] That is, we abstract from the intrinsic dynamic consequences of adopting a project—e.g. affecting the remaining pool of potential projects, or affecting the needs/preferences of the agent going forward. In this sense, we isolate the dynamic delegation problem.

[2] This standard solution is sometimes called "selling the firm." For instance, see p. 482 of Mas-Colell, Whinston, and Green (1995).

[3] This assumption is stronger than needed. As long as the agent cannot make transfers to the principal, our main results are qualitatively unchanged.

Defense , for example, is unlikely to ask soldiers to pay for their own body armor.

Our first result, Theorem 1, is an efficiency bound on any delegation rule that involves only good projects being taken. If no bad projects are initiated, the principal and the agent have aligned interests, both preferring more projects. The best such delegation rule, the **Aligned Optimal Budget**, has a very simple form. The principal delegates to the agent until the agent adopts a project, but she follows any project with a temporary freeze. That is, no more projects are allowed for the next $\bar{\tau}$ units of time. Then, the same contract starts over. The optimal $\bar{\tau}$ will be the shortest freeze duration severe enough to keep the agent from taking bad projects. In the context of a university, the physics department can freely search for a candidate, but any hire is followed by a temporary hiring freeze. During the freeze, although many qualified candidates may be available, the department is forbidden from hiring them. We show that the resulting inefficiency remains even if the parties are arbitrarily patient.

Our main result, Theorem 2, is a full characterization of the optimal intertemporal delegation rule. The uniquely[4] optimal contract, the **Dynamic Capital Budget**, comprises two distinct regimes. At any time, the parties engage in either Capped Budgeting or Controlled Budgeting.

In the **Capped Budget** regime, the principal always delegates, and the agent initiates all good projects that arrive. At the relationship's outset, the agent has an expense account for projects, indexed by an initial balance and an account balance cap. The balance captures the number of projects that the agent could adopt immediately without consulting the principal. Any time the agent takes a project, his balance declines by 1. While the agent has any funds in his account, the account accrues interest. If the agent takes few enough projects, the account will grow to its cap. At this balance, the agent is still allowed to take projects, but his account doesn't grow any larger (even if he waits). Not being rewarded for fiscal restraint, the agent immediately initiates a project, and his balance again declines by 1.

If the agent overspends, a **Controlled Budget** regime begins: the principal first imposes a temporary freeze to punish the agent: a larger overdraft is met with a

[4]More precisely, the two-regime structure and the exact character of the Capped Budget regime are uniquely required by optimality.

longer freeze. The parties then revert to the Aligned Optimal Budget. It is worth noting that the players are certain to eventually enter this regime. Once there, the Controlled Budget regime is absorbing.

With a Capped Budget, the principal tolerates some bad projects, and, in turn, avoids freezes. For a well-chosen account balance, the efficiency gain (more search time for good projects) outweighs the risk of bad projects.

There are two broader lessons to be learned from the above characterization. First, the disadvantaged principal, stripped of all her usual tools, can nonetheless leverage the agent's information with some success. Second, the optimal equilibrium exhibits rich dynamics for such a simple, stationary model.

Prima facie, the only incentivizing instrument available to the principal is mutual "money burning" in the form of (temporarily) freezing. For instance, a freeze on equipment acquisition by the Department of Defense can be a useful threat, inducing frugal decisions now, but it comes at a cost: its own soldiers will sometimes be unequipped even in times of need. This force, with the cost it entails, single-handedly disciplines the agent under Controlled Budgeting. However, the principal has an additional tool: the expectation of future lenience can serve as a reward for the agent today. To induce frugal decisions now, the Department of Defense may promise more budgetary freedom in the future. The Capped Budget makes use of both this reward and above punishment—*carrot* and *stick*. A high account balance entails the promise of future permissiveness from the principal, while a low account balance entails an imminent threat of Controlled Budgeting. When the budget is below the cap, the principal rewards the agent for his diligence with the account's interest accrual. As long as the promise is credible—i.e. the principal would rather fulfill her contract than unilaterally freeze the relationship—the reward will be credible too. At the cap itself, the principal cannot credibly promise further lenience, and good behavior by the agent would go unpaid; accordingly, the agent takes a project immediately. If the unit has shown enough fiscal restraint, the Department of Defense purchases new equipment, independent of its need, to reward the unit.

The optimal contract yields clear dynamics for the delegation relationship. Both regimes reflect a productive relationship, but each is of a distinct character. Capped

Budgeting is highly productive but low-yield:[5] every good project is adopted, but some bad projects are as well. Controlled Budgeting is high-yield but less productive: only good projects are adopted, but some good opportunities go unrealized. In this sense, as the Capped Budget regime is transient, the relationship naturally drifts toward conservatism. The principal's payoff comparisons among regimes are ambiguous: at lower budget balances, Capped Budgeting dominates Controlled Budgeting, while the relationship reverses for larger balances.

The remainder of the paper is structured as follows. In the following pages, we discuss the related literature. Section 2 presents the model and introduces a convenient language for discussing players' incentives in our model. In Section 3, we discuss aligned equilibria—i.e. those in which no bad projects are adopted; we characterize the class and show that such equilibria are necessarily inefficient. The heart of the paper is Section 4, in which we present the Dynamic Capital Budget contract and prove its optimality. In Section 5, we discuss some possible extensions of our model. Final remarks follow in Section 6.

Related Literature

This paper belongs to a rich literature on delegated decision making,[6] initiated by Holmström (1984), wherein a principal faces a tradeoff between leveraging an agent's private information and shielding herself from his conflicting interests. The key issue is how much freedom the principal should give the agent; the more aligned their preferences are, the more discretion she should allow. Armstrong and Vickers (2010) find that the principal optimally excludes some ex-post favorable options in order to provide better incentives to the agent ex-ante, while Ambrus and Egorov (2013) highlight the value created by money burning as a means to alleviate incentive constraints. These insights apply to our model, in which indirect money burning—harmful to both players—is used to provide incentives.

Our paper contributes to the recently active field of dynamic delegation. Malenko (2013) characterizes the optimal contract for a principal who delegates investment

[5] By "productive," we mean that a lot of value is delivered to the agent. By "low-yield," we mean that less such value is delivered per unit of cost to the principal.

[6] For instance, see Frankel (2014) and the thorough review therein.

choices and has a costly state verification technology, monetary transfers, and commitment power: a capital expense account with a fluctuating interest rate. Guo (2014) focuses on the delegation of a dynamic decision problem to an agent with non-transient private information. Alonso and Matouschek (2007) indicate how dynamic threats can partially bridge the gap between the cheap-talk and delegation models. In contemporaneous work, Guo and Hörner (2014) study optimal dynamic mechanisms without money in a world of partially persistent valuations, in which the principal has commitment power. The principal's ability to commit generates different incentive dynamics: in contrast to our model, the agent may receive his first-best outcome in the long run.

Our model speaks to the relational contracting literature, as in Pearce and Stacchetti (1998), Levin (2003), and Malcomson (2010). This literature focuses on relationships in which formal contracting is impossible, and all incentives—and the credibility of promises that provide those incentives—are anchored to the future value of the relationship. In particular, Li and Matouschek (2013) focus on the case in which the principal's opportunity cost of promise keeping is private information. In both their model and ours, the inability to formally contract leads a stationary problem to be met with a non-stationary relationship. In theirs, the relationship is cyclical, with every punishment being strictly temporary. In ours, the relationship temporarily cycles, before drifting toward conservatism.

Our results add to the literature on relationship building under private information. One strand of the literature concerns itself with the building and maintenance of partnerships, such as Möbius (2001), Hauser and Hopenhayn (2008), and Espino, Kozlowski, and Sanchez (2013). Möbius (2001) constructs a model in which players privately observe opportunities to do favors for one another at personal cost; Hauser and Hopenhayn (2008) indicate that the relationship can benefit from varying incentives based on both action and inaction. In a related strand of the literature, Chassang (2010) and Li, Matouschek, and Powell (2015) focus on the relationship between a firm and its employee, whose private information can generate persistent differences in performance across ex-ante identical firms. In contemporaneous work, Li, Matouschek, and Powell (2015) focus on a repeated trust game—the principal either takes a safe option or trusts the biased, but better-informed, agent—

preceded at every stage by a simultaneous entry decision. If either player chooses not to enter the game, both are uniformly punished. The opportunity to unilaterally punish the principal makes long-term reward for the agent credible, generating history dependence similar to that in Guo and Hörner (2014): owing to the firm's inability to interpret its employee's actions, the realization of random early outcomes has long-lasting consequences. A final strand of this literature regards dynamic corporate finance, as in Clementi and Hopenhayn (2006) or Biais et al. (2010). In Biais et al. (2010), the principal commits to investment choices and monetary transfers to the agent, who privately acts to reduce the chance of large losses for the firm. While our setting is considerably different, their optimal contract and ours exhibit similar dynamics: our "funny money" balance takes the role of real sunk investment.

Lastly, there is a deep connection between the present work and the literature on linked decisions. Casella (2005) and Jackson and Sonnenschein (2007) consider a setting in which, given a large number of physically independent decisions, the ability to connect them across time helps align incentives. Frankel (2011, 2013) considers environments in which a principal with commitment power optimally employs a budgeting rule to discipline the agent. Linking decisions across time, as in the above literature, is always possible if the principal can commit to a budgetary rule. In our model—without such commitment power—dynamic budgeting remains optimal but is tempered by the principal's need for credibility.

2 The Model

We consider an infinite-horizon two-player ($\mathcal{P}$rincipal and $\mathcal{A}$gent) game in discrete time. Each period, the principal chooses whether or not to delegate a project adoption choice to the agent. Conditional on delegation, the agent privately observes which type of project is available and then publicly decides whether or not to adopt it. At the time of its adoption, a project of type θ generates an agent payoff of θ. Each project entails an implementation cost of c, to be borne solely by the principal; thus, a project yields a net utility of $\theta - c$ to the principal. Notice that the cost is independent of the project's type. In particular, the difference between the agent's payoffs and the principal's payoffs doesn't depend on the agent's private information. We interpret this payoff structure as the principal innately caring about the agent's (unobservable) payoff, in addition to the cost that she alone bears. While the university's president cannot expertly assess a specialized candidate, she still wants the physics department to hire good physicists. The state government can't assess the added value of each local public project, but it still values the benefit that a project brings to the community. Given this altruistic motive, the principal cares about the value generated by a project, even though she never observes it.

While the players rank projects in the same way, the key tension in our model is a disagreement over which projects are worth taking. The agent cares only about the benefit generated by a project, while the principal cares about said benefit net of cost; we find revenue and profit to be useful interpretations of the players' payoffs.

$\mathcal{P}$ and $\mathcal{A}$ share a common discount factor $\delta \in (0, 1)$, maximizing *expected discounted profit* and *expected discounted revenue*, respectively. So, if the available project in each period $t \in \mathbb{Z}_+$ is θ_t and projects are adopted in periods $\mathcal{T} \subseteq \mathbb{Z}_+$, then the principal and agent get profit and revenue,

$$\Pi = \sum_{t \in \mathcal{T}} \delta^t(\theta_t - c) \text{ and } V = \sum_{t \in \mathcal{T}} \delta^t \theta_t, \text{ respectively.} \tag{1}$$

First, $\mathcal{P}$ publicly decides whether to *freeze* project adoption or to *delegate* it. If $\mathcal{P}$ freezes, nothing happens and both players accrue no payoffs. If $\mathcal{P}$ delegates, $\mathcal{A}$ privately observes which type of project is available and decides whether or not to initiate the available project. The current period's project is good (i.e. of type $\bar{\theta}$) with probability $h \in (0, 1)$ and bad (i.e. of type $\underline{\theta}$) with complementary probability. If the agent initiates a project of type θ, payoffs $(\theta - c, \theta)$ accrue to the players. The principal observes whether or not the agent initiated a project, but she never sees the project's type.

Notation. *Let $\theta_E := (1 - h)\underline{\theta} + h\bar{\theta}$ be the* ***ex-ante expected project value****.*

Throughout the paper, we maintain the following assumption:

Assumption 1.

$$0 < \underline{\theta} < \theta_E < c < \bar{\theta}.$$

Each period, $\mathcal{P}$ and $\mathcal{A}$ play the following stage game:

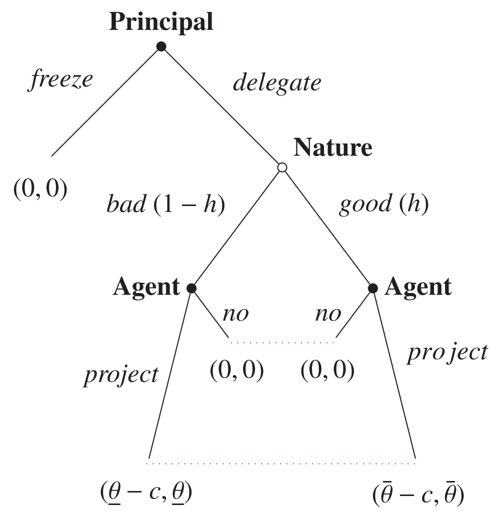

Figure 1: The principal observes the agent's choices but not project quality.

Assumption 1 characterizes the preference misalignment between agent and principal. Since $\underline{\theta} - c < 0 < \bar{\theta} - c$, the principal prefers good projects to nothing, but prefers inactivity to bad projects. Given $0 < \underline{\theta} < \bar{\theta}$, the agent prefers any project to no project, but also prefers good ones to bad ones. So, they agree on which projects are better to adopt, but may disagree on whether a given project is worth taking ex-post. The condition $\theta_E < c$ (interpreted as an assumption that good projects are scarce) says that the latter effect dominates, and the conflict of interest prevails even ex-ante: the principal prefers a freeze to the average project. A good enough physicist is rare; the university finds hiring worthwhile only if it can rely on the department to separate the wheat from the chaff. If the players interacted only once, the department would not be selective. Accordingly, the stage game has a unique sequential equilbrium: the principal freezes, and the agent takes a project if allowed.

Equilibrium Values

Throughout the paper, *equilibrium* will be taken to mean perfect semi-public equilibrium (PPE), in which the players respond only to the public history of actions and (for the agent) current project availability. While in a different setting, this definition in similar in spirit to that in Compte (1998) and in Harrington and Skrzypacz (2011).

Definition 1. *Each period, one of three public outcomes occurs: the principal freezes; the principal delegates and the agent initiates no project; or the principal delegates and the agent initiates a project. A time-t* ***public history****, h_t, is a sequence of t public outcomes (along with realizations of public signals). The agent has more relevant information when making a decision. A time-t* ***agent semi-public history*** *is $h_t^{\mathcal{A}} = (h_t, D, \theta_t)$, where h_t is a public history, D is a principal decision to delegate, and θ_t is a current project type.*

A ***principal public strategy*** *specifies, for each public history, an action: delegate or freeze. An* ***agent semi-public strategy*** *specifies, for each agent semi-public history, an action: project adoption or no project adoption.*

A ***perfect (semi-)public equilibrium*** *(PPE) is a sequential equilibrium in which the principal plays a public strategy, while the agent plays a semi-public strategy.*

Every equilibrium entails an expected discounted number of adopted good projects $g = \mathbb{E} \sum_{t \in \mathcal{T}} \delta^t \mathbf{1}_{\{\theta_t = \bar{\theta}\}}$ and an expected discounted number of adopted bad projects $b = \mathbb{E} \sum_{t \in \mathcal{T}} \delta^t \mathbf{1}_{\{\theta_t = \underline{\theta}\}}$, where $\mathcal{T} \subseteq \mathbb{Z}_+$ is the realized set of periods in which the principal delegates and the agent adopts a project. Given those, one can compute the agent value/revenue as

$$v = \bar{\theta} g + \underline{\theta} b$$

and the principal value/profit as

$$\pi = (\bar{\theta} - c)g - (c - \underline{\theta})b.$$

For ease of bookkeeping, it is convenient to track equilibrium-supported revenue v and bad projects b, both in expected discounted terms. The vector (v, b) encodes both agent value v and principal profit

$$\begin{aligned}\pi(v,b) &= (\bar{\theta}-c)g-(c-\underline{\theta})b \\ &= (\bar{\theta}-c)\frac{v-\underline{\theta}b}{\bar{\theta}}-(c-\underline{\theta})b \\ &= \left(1-\frac{c}{\bar{\theta}}\right)v-c\left(1-\frac{\underline{\theta}}{\bar{\theta}}\right)b.\end{aligned}$$

Toward a Characterization The main objective of this paper is to characterize the set of equilibrium-supported payoffs,

$$\mathcal{E}^* := \{(v,b) : \exists \text{ equilibrium with revenue } v \text{ and bad projects } b\} \subseteq \mathbb{R}^2_+.$$

Throughout the paper, we make extensive use of two simple observations about the set $\mathcal{E}^*$. First, notice that $(0,0) \in \mathcal{E}^*$, since the profile σ^{static}, in which the principal always freezes and the agent takes every permitted project, is an equilibrium. Said differently, there is always an unproductive equilibrium—i.e. one with no projects. That this equilibrium provides min-max payoffs makes our characterization easier. Second, as the following lemma clarifies, off-path strategy specification is unnecessary in our model. For any profile satisfying appropriate on-path incentive constraints, one can always find another profile with identical on-path behavior, but altered off-path to make the profile an equilibrium. With the lemma in hand, we rarely specify off-path behavior in a given strategy profile, as we are chiefly interested in payoffs.

Lemma 1. *Fix a strategy profile σ, and suppose that:*

1. *The agent has no profitable deviation from any on-path history.*
2. *Following all on-path histories, the principal has nonnegative continuation profit.*

Then, there is an equilibrium $\tilde{\sigma}$ that generates the same on-path behavior (and, therefore, the same value profile).

Proof. Let σ^{static} be the stage Nash profile—i.e. the principal always freezes, and the agent takes a project immediately whenever permitted.

Define $\tilde{\sigma}$ as follows.

- On-path (i.e. if $\mathcal{P}$ has never deviated from σ's prescription), play according to σ.
- Off-path (i.e. if $\mathcal{P}$ has ever deviated from σ's prescription), play according to σ^{static}.

The new profile is incentive-compatible for the agent: off-path because σ^{static} is, on-path because σ is. It is also incentive-compatible for the principal: off-path because σ^{static} is, on-path because σ is and has nonnegative continuation profits while σ^{static} yields zero profit. □

Dynamic Incentives

While our results concern a discrete-time repeated game, we find it expositionally convenient to present the intuition in continuous time. We present results for the case in which the players interact very frequently, but good projects remain scarce. A unit can find desirable equipment to request from the Department of Defense at any time; what is rare is the opportunity to buy equipment whose benefit offsets its cost. The cleanest economic intuition lies in this limiting case. Letting the time between decisions, together with the proportion of good projects, vanish enables us to present our main results heuristically in the language of calculus.

Given a period length $\Delta > 0$, we follow a standard parametrization: discount factor $\delta = 1 - r\Delta$, and proportion of good projects $h = \eta\Delta$ for fixed $r, \eta > 0$. In the limit, as $\Delta \to 0$, good projects arrive with Poisson rate η, and bad projects are always available. Rather than yielding flow payoffs, an initiated project of type θ provides the players a lump-sum revenue of θ, at a lump-sum cost[7] (to the prin-

[7] The analysis would not be changed if the benefit had a flow component but the cost were lump-sum, in which case θ would be interpreted as a present discounted value. If the cost were not lump-sum, on the other hand, the principal would face a new incentive constraint—to willingly continue to fund a costly project.

cipal) of c. A park funded by the state government, for instance, serves the local community but at a sizable construction cost to the state.

Throughout the analysis, we normalize $r = 1$ and interpret η as the ratio $\frac{\eta}{r}$, the good project arrival rate per effective unit of time. While we focus on the limit as $\Delta \to 0$ and thus $\delta \to 1$, the present work is not a folk theorem analysis.[8] Finally, observe that in the limiting case $\theta_E = \underline{\theta}$.

Self-Generation If we aim to understand the players' incentives at any given moment, we must first understand how their future payoffs evolve in response to their current choices. To describe the law of motion of revenue v [or, respectively, bad projects b], we keep track of:

- $\dot{v}_t$ [resp. $\dot{b}_t$], the rate of change of v [resp. b] conditional on no project adoption; and
- $\tilde{v}_t$ [resp. $\tilde{b}_t$], the continuation of v [resp. b] if a project is undertaken.

Observe that the continuation values cannot depend on the quality of the adopted project (nor can the laws of motion depend on availability of forgone projects), which is not publicly observable. Finally, describe the players' present actions as follows:

- The principal makes a delegation choice[9] $d \in \{0, 1\}$,—i.e. whether or not to let the agent execute a project in the current period.
- The agent chooses $\hat{\eta} \in [0, \eta]$ and $\hat{\lambda} \in [0, \infty]$, the instantaneous rates at which he currently initiates good and bad projects, respectively, conditional on being allowed to.

[8] A folk theorem analysis would entail taking $r \to 0$ for a fixed arrival rate of good projects, and thus taking their ratio $\eta \to \infty$. The distinction is analogous to that in Abreu, Milgrom, and Pearce (1991).

[9] As we show in the appendix, it is without loss of generality that the principal uses a pure strategy. The intuition is that, since everything $\mathcal{P}$ does is publicly observed, any mixing she does may as well be replaced with public mixing.

By way of interpretation, $d_t, \hat{\eta}_t, \hat{\lambda}_t$ are the approximate choices on either $[t, t+\Delta)$ for small $\Delta > 0$ or until the next project, whichever comes first.[10] In particular, $d_t = 1$ doesn't mean that the agent has the opportunity to initiate an arbitrarily large number of bad projects without the principal's continued consent.

Appealing to self-generation arguments, as in Abreu, Pearce, and Stacchetti (1990), and to Lemma 1, equilibrium is characterized by the following three conditions:

1. Promise keeping:[11]

$$
\begin{aligned}
v &= d\left[(\hat{\eta}\bar{\theta} + \hat{\lambda}\underline{\theta}) - (\hat{\eta} + \hat{\lambda})(v - \tilde{v})\right] + \dot{v} \\
&= d\hat{\eta}(\bar{\theta} + \tilde{v} - v) + d\hat{\lambda}(\underline{\theta} + \tilde{v} - v) + \dot{v} \\
b &= d\left[\hat{\lambda} - (\hat{\eta} + \hat{\lambda})(b - \tilde{b})\right] + \dot{b} \\
&= d\hat{\eta}(0 + \tilde{b} - b) + d\hat{\lambda}(1 + \tilde{b} - b) + \dot{b}.
\end{aligned}
$$

   We decompose continuation outcomes (v, b) from any instant into what happens in each of three events—the agent finds and invests in a good project at that instant; the agent finds and invests in a bad project at that instant; no project is adopted—weighted by their instantaneous probabilities.

2. Agent incentive compatibility:

$$
v - \tilde{v} \begin{cases} \geq \underline{\theta} & \text{if } \hat{\lambda} < \infty \\ \leq \bar{\theta} & \text{if } \hat{\eta} > 0. \end{cases}
$$

   If the agent is willing to resist taking a project immediately ($\hat{\lambda} < \infty$), it must

[10] Notice that, if the principal chooses $d_t = 1$ and the agent chooses $\hat{\lambda}_t = \infty$, then both players face new choices to make, still exactly at time t.

[11] When $d = 1$ and $\hat{\lambda} = \infty$, we replace the given equations with the limiting equations obtained from dividing through by $\hat{\lambda}$:

$$
\begin{aligned}
0 &= \underline{\theta} - (v - \tilde{v}) \\
0 &= 1 - (b - \tilde{b}).
\end{aligned}
$$

When $d = 0$ and $\hat{\lambda} = \infty$, we let $d\hat{\lambda} = 0$, so that the principal retains ultimate authority over project adoption.

be that the punishment $v - \tilde{v}$ for taking a project is severe enough to deter the $\underline{\theta}$ myopic gain; similarly, if the agent is to take some good projects ($\hat{\eta} > 0$), the same punishment $v - \tilde{v}$ cannot be too draconian.

3. Principal participation:
$$\pi(v, b) \geq 0.$$
The principal could, at any moment, unilaterally move to a permanent freeze and secure herself a profit of zero. Therefore, at any history, she must be securing at least that much in equilibrium.

3 Aligned Equilibrium

We have established that our game has no productive stationary equilibrium. If the principal allows history-independent project adoption, the agent cannot be stopped from taking limitless bad projects. In the present section, we ask whether this core tension can be resolved by allowing non-stationary equilibria. More precisely, are there productive aligned equilibria?

Definition. *An* ***aligned equilibrium*** *is an equilibrium in which no bad projects are ever adopted.*

A sensible first attempt is to delegate, but to punish the agent as much as possible as soon as he might have taken a bad project. Describe σ^∞ as follows: the principal allows exactly one project, after which she shuts down forever; the agent takes the first good project that comes along. Is σ^∞ an equilibrium? For this profile, before the first project,
$$d = 1, \;\; \hat{\eta} = \eta, \;\; \hat{\lambda} = 0, \;\; \tilde{v} = 0, \text{ and } \dot{v} = 0.$$
Therefore, promise keeping gives
$$v = \eta(\bar{\theta} + 0 - v) + 0 \implies v = \frac{\eta}{1+\eta}\bar{\theta} \;\; \in (0, \bar{\theta}).$$
Principal participation is immediate when there are no bad projects, so that we only

need to check agent incentive compatibility, which holds if and only if

$$v - \tilde{v} \geq \underline{\theta} \iff \frac{\eta}{1+\eta}\bar{\theta} \geq \underline{\theta} \iff \eta(\bar{\theta} - \underline{\theta}) \geq \underline{\theta}.$$

Notation. *Let $\omega := \eta(\bar{\theta} - \underline{\theta})$ be the **marginal value of search**.*
The constant ω captures the marginal option value of searching for a good project until the next instant.

Assumption 2.

$$\omega > \underline{\theta}.$$

Unless otherwise stated, we will assume that Assumption 2 holds throughout. In discrete time, Assumption 2 can equivalently be expressed as a lower bound on the discount factor δ. If the agent is sufficiently patient, the marginal value of searching for a good project outweighs the myopic benefit of an immediate bad project.

A Stick with No Carrot

The argument above demonstrates that the threat of shutdown is enough to incentivize picky project adoption. However, permanent shutdown destroys a lot of value, for both the principal and the agent. If the university allows the physics department only one hire for its entire existence, every good candidate is passed over thereafter, harming the university. It is natural to ask whether a less severe mutual punishment can provide the same incentives.

Given $\tau \in (0, \infty]$, describe the τ-freeze stationary contract σ^τ as follows:

1. The principal starts by delegating, and does so indefinitely if no projects are taken.

2. The agent takes no bad projects, and takes the first good project that arrives.

3. Any project is followed by a freeze of length τ, followed by restarting σ^τ.

We can interpret the τ-freeze contract as a simple budget rule. The agent is given a budget of one project by the principal. If $\mathcal{A}$ does not spend his budget, it rolls over to the next instant. If the budget is depleted, $\mathcal{P}$ replenishes it after a waiting period τ. The physics department can take as much time as needed to find a suitable candidate, but the hire is followed by a two-year freeze; afterward, the university allows the department to search again.

As seen in the previous section, σ^∞ is an equilibrium. By continuity, σ^τ is an equilibrium for sufficiently high finite τ. Moreover, for $\tau \approx 0$, the contract σ^τ cannot be an equilibrium. Indeed, in a delegation phase,

$$v - \tilde{v} = (1 - e^{-\tau})v \leq (1 - e^{-\tau})\eta\bar{\theta} \xrightarrow{\tau \to 0} 0.$$

In particular, the punishment for executing a project, $v-\tilde{v}$, is smaller than the benefit of an immediate project, $\underline{\theta}$, for sufficiently small τ. Hence, the agent strictly prefers to take the (bad) project in front of him.

The revenue generated by σ^τ is decreasing in τ: less shutdown means fewer forgone opportunities for good projects, which means more revenue. What is less clear is how the punishment $v - \tilde{v}$ changes with τ. As we increase τ, the punishment is $v - \tilde{v} = (1 - e^{-\tau})v$, which increases as a fraction of total revenue. Thus, its comparative statics are not obvious, as increasing τ makes the punishment a bigger share of a smaller pie. Let's compute it: promise keeping gives $v = \eta[\bar{\theta} + \tilde{v} - v] + 0 = \eta[\bar{\theta} - (1 - e^{-\tau})v]$, which implies

$$v = \frac{\eta}{1 + \eta(1 - e^{-\tau})}\bar{\theta} \text{ and } v - \tilde{v} = \frac{\eta(1 - e^{-\tau})}{1 + \eta(1 - e^{-\tau})}\bar{\theta}.$$

With the revenue decreasing in τ and the punishment increasing in τ, the following proposition follows readily.

Proposition 1. *Consider* $\{\sigma^\tau\}_{\tau\in(0,\infty]}$ *as above.*

1. *There is a unique* $\bar{\tau} \in (0, \infty]$ *satisfying* $\dfrac{\eta(1 - e^{-\bar{\tau}})}{1 + \eta(1 - e^{-\bar{\tau}})}\bar{\theta} = \underline{\theta}$.
2. σ^τ *is an equilibrium if and only if* $\tau \geq \bar{\tau}$.

3. *Among all such τ, the choice $\bar{\tau}$ provides the highest revenue (and, thus, the highest profit).*

4. *The revenue of $\sigma^{\bar{\tau}}$ is ω, and so its profit is $(1 - \frac{c}{\theta})\omega$.*

Proof. Everything is proven above, except for the expression for $\bar{\tau}$ and the generated revenue. To compute $\bar{\tau}$,

$$\begin{aligned}
\frac{\eta(1-e^{-\tau})}{1+\eta(1-e^{-\tau})}\bar{\theta} = \underline{\theta} \quad &\Longleftrightarrow \quad \eta(1-e^{-\tau})\bar{\theta} = [1+\eta(1-e^{-\tau})]\underline{\theta} \\
&\Longleftrightarrow \quad (1-e^{-\tau})\omega = \underline{\theta} \\
&\Longleftrightarrow \quad \tau = \log\frac{\omega}{\omega-\underline{\theta}}.
\end{aligned}$$

The associated revenue is then

$$\begin{aligned}
v &= \frac{\eta}{1+\eta(1-e^{-\bar{\tau}})}\bar{\theta} \\
&= \frac{\eta\bar{\theta}}{1+\eta\frac{\underline{\theta}}{\omega}} = \frac{\eta\bar{\theta}}{\omega+\eta\underline{\theta}}\omega \\
&= \omega.
\end{aligned}$$

□

This simple class of contracts illuminates the forces at play in our model. The principal wants good projects to be initiated, but she cannot afford to give the agent free rein. If she wants to stop him from investing in bad projects, she must threaten him with mutual money burning. Subject to wielding a large enough stick to encourage good behavior, she efficiently wastes as little opportunity as possible. The university does not want to deprive the physics department of needed faculty, and so it should limit them only enough to discipline fiscal restraint.

One may be concerned that frequent shutdown leaves a lot of opportunities unrealized. Accordingly, it seems sensible to seek other plausible aligned equilibria in which less value is destroyed.

A Smaller Stick

What if, instead of allowing one project followed by temporary freeze, $\mathcal{P}$ allows $K \in \mathbb{N}$ projects, before freezing for $\tau \in (0, \infty]$? The main lesson in Jackson and Sonnenschein (2007) is that budgetary rationing of multiple decisions can alleviate incentive misalignment, at a minor welfare cost. One might, therefore, hope that allowing K projects before punishing enables a more productive relationship, while still incentivizing the agent to avoid bad projects. As it turns out, this affords no real improvement. If the (now more distant) punishment for the first project is to be enough to stop the agent from cheating initially, it must be severe enough to negate the would-be benefits of delayed closure. With computations similar to those in the previous section, it is straightforward that: the "K-project, τ-freeze" strategy profile is an equilibrium if and only if it delivers an initial value $\leq \omega$ to the agent. So the principal can allow more projects before punishing the agent (and herself), but if she is to still deter the agent from cheating, she has to make the punishment phase—which happens farther in the future—longer for bigger K. This makes higher K redundant: no such equilibrium can outperform the "1-project, $\bar{\tau}$-freeze" equilibrium.

Aligned Optimality

The preceding analysis suggests a fundamental limit to how productive an aligned equilibrium can be. Indeed, with no bad projects, the principal has only one dimension—expected discounted good projects or, equivalently, $\mathbb{E}\int e^{-t}\eta\mathbf{1}_{\{\text{open at time } t\}}\,\mathrm{d}t$—with which to provide incentives. Delaying a punishment, making it less likely, or making it less severe are all different physical instruments to alleviate the same money-burning cost, but with the same adverse effect on agent incentives. The following theorem shows that the upper bound that we have uncovered in specific classes of equilibria—the marginal value of search—is no coincidence.

Theorem 1 (Aligned Optimality)**.** [12]

1. *There exist productive aligned equilibria if and only if*[13] *Assumption 2 holds.*

[12]The discrete time counterpart is Theorem 1, in the appendix.

[13]We abstract from the knife-edge case $\omega = \underline{\theta}$.

2. *Every aligned equilibrium generates revenue less than or equal to the marginal value of search, ω.*

3. *The $\bar{\tau}$-freeze contract, where $\bar{\tau} = \log \frac{\omega}{\omega - \underline{\theta}}$, is, therefore, optimal among all aligned equilibria, given Assumption 2.*

Proof. Following any history, in any aligned equilibrium,

$$
\begin{aligned}
\dot{v} &= v - p\left[(\hat{\eta}\bar{\theta} + \hat{\lambda}\underline{\theta}) - (\hat{\eta} + \hat{\lambda})(v - \tilde{v})\right] \\
&= v - p\left[(\hat{\eta}\bar{\theta} + 0\underline{\theta}) - (\hat{\eta} + 0)(v - \tilde{v})\right] \\
&= v - p\hat{\eta}[\bar{\theta} - (v - \tilde{v})] \\
&\geq v - p\hat{\eta}[\bar{\theta} - \underline{\theta}] \text{ (since agent IC \& no bad projects } \implies v - \tilde{v} \geq \underline{\theta} \text{ if } p > 0) \\
&\geq v - \eta(\bar{\theta} - \underline{\theta}) \\
&= v - \omega.
\end{aligned}
$$

So, if $\epsilon := v_0 - \omega > 0$, then v grows indefinitely at rate $\dot{v} \geq v - \omega \geq v_0 - \omega = \epsilon$, so that $v_t \geq v_0 + t\epsilon$. This would contradict the fact that $(v_t, b_t) \in \bar{\mathcal{E}}$ (a compact set) for every history. Thus, it must be that $v_0 \leq \omega$, verifying (2).

For (1), suppose that Assumption 2 is violated. Consider any productive equilibrium σ. Dropping to an on-path history if necessary, we may assume that σ doesn't start with a freeze. If $\hat{\lambda}_0 > 0$, then σ isn't an aligned equilibrium. If $\hat{\lambda}_0 < \infty$, then agent IC implies $v_0 - \tilde{v}_0 \geq \underline{\theta}$. Therefore,

$$v_0 \geq v_0 - \tilde{v}_0 \geq \underline{\theta} > \omega.$$

Appealing to the first part, it must be that σ is not an aligned equilibrium.

Under Assumption 2, $\sigma^{\bar{\tau}}$ is an equilibrium providing revenue ω, proving (3) and the remaining direction of (1). □

The second result above gives a firm upper bound on how much value can be created in an aligned equilibrium. If the principal wants the agent to behave, she has to stop him from taking bad projects. In an aligned equilibrium—in which the principal's payoffs are directly proportional to the agent's—anything that punishes the agent punishes the principal just as much. Since the rolling budget rule $\sigma^{\bar{\tau}}$ entails as little punishment as possible subject to agent IC whenever the principal is delegating, it is best for both players within the class of aligned equilibria. In what follows, we refer to $\sigma^{\bar{\tau}}$ as our **Aligned Optimal Budget**.

One important consequence of the above is that aligned equilibria cannot hope to achieve first-best for the principal, even as the players become very patient.[14]

Corollary 1. *The ratio of aligned optimal profit to first-best profit is*

$$\frac{\omega(1 - \frac{c}{\bar{\theta}})}{\eta(\bar{\theta} - c)} = \frac{\bar{\theta} - \underline{\theta}}{\bar{\theta}} < 1,$$

which is independent of players' patience.

Expiring Budget: Beyond Aligned Equilibria

We now consider an intuitive class of contracts, showcasing a new incentivizing tool available to the principal. In aligned equilibria, punishment via mutually costly shutdown bears the full cost of providing incentives. In addition to such punishment, the principal has a means to reward the agent: a project—no questions asked. The principal can delegate to the agent t periods in the future, conditioning no other decisions on the agent's choice. In doing so, she gifts $e^{-t}\theta_E$ to the agent, at a personal cost of $e^{-t}(c - \theta_E)$. As multiple such projects can be awarded at various times, this mechanism amounts to transferable utility. Perhaps, by allowing bad projects in case the agent has enough bad luck, the principal can benefit from burning less value following good luck.

Consider the Expiring Budget contract, in which $\mathcal{A}$ is allowed to adopt one project per calendar year—use it or lose it. $\mathcal{P}$ delegates until $\mathcal{A}$ takes a project, at which point $\mathcal{P}$ freezes until the end of the year. If the agent doesn't use his one-project budget during that year, it expires. Each fiscal year, $\mathcal{A}$ takes the first good project that arrives but resorts to a bad one if, by the end of the year, no good project has arrived.

[14]As discussed in Footnote 8, we can understand the patient limit as $\eta \to \infty$

Under Assumption 2, one can verify that the the agent optimally exerts restraint during the year; and at the last instant of a fiscal year, facing no opportunity cost, the agent spends his imminently expiring budget. The presence of bad projects introduces a new concern for equilibrium: $\mathcal{P}$'s credibility. As the end of the year approaches, if the agent has not yet used his annual budget, the likelihood of any good project arriving this year vanishes. Equilibrium requires that, when called to deliver an immediate (likely bad) project at year's end, the principal would rather do so than sever the relationship.

While it may be profit-enhancing to reward the agent with bad projects following fiscal restraint,[15] equilibrium requires that the promise of such bad projects be credible.

4 Dynamic Capital Budgeting

In the previous section, we presented some sensible budget rules. First, in aligned equilibria, the agent is punished via mutual money burning for each project but cannot be rewarded for fiscal restraint. Next, in Expiring Budget contracts, the agent is both punished (shutdown) for taking a project and rewarded (a project, no questions asked) for waiting. The form of this reward, however, is inefficient. Conditional on reaching the first fiscal year's end with no good project, the agent is certain to adopt a bad one. What if, instead of forcing an unused project to expire, the principal adds it to next year's budget? As we learned in the "A Smaller Stick" subsection, an agent with a larger budget is less diligent because the search opportunity cost of a project is smaller. An impatient agent would still initiate one project immediately, leaving a budget of one project for the second year; nothing changes relative to the Expiring Budget rule. A more patient agent, however, would save his "stock" of projects for the future and, in doing so, make both players better off. The agent spends the second year searching for up to two good projects, and only at year's end (if his bad luck continues) liquidates his budget. For this modified budget to be an equilibrium, the principal's promise must remain credible, even if

[15] We show in the appendix that, for some parameter values, an Expiring Budget outperforms any aligned equilibrium.

the agent's search is unsuccessful. At the end of the second year, the principal must prefer to finance two immediate projects (again, most likely bad) rather than sever the relationship.

There are efficiency gains to be had from smoothing the agent's budget across time, but one must carefully balance the principal's credibility constraint for it to remain an equilibrium. Our next rule, the **Dynamic Capital Budgeting** (DCB) contract, is an attempt to achieve this balance.

The DCB contract is characterized by a budget cap $\bar{x} \geq 0$ and initial budget balance $x \in [-1, \bar{x}]$, and consists of two regimes. At any time, players follow Controlled Budgeting or Capped Budgeting, depending on the agent's balance, x. The account balance can be understood as the number of projects the agent can initiate without immediately affecting the principal's delegation decisions.

Capped Budget $(x > 0)$

The account balance grows at the interest rate $r = 1$ as long as $x < \bar{x}$. Accrued interest is used to reward the agent for fiscal restraint. Since the search opportunity cost of taking a project decreases in the account balance, the reward for diligence is increasing (exponentially) to maintain incentives. While $\mathcal{A}$'s account is in the black, $\mathcal{P}$ fully delegates project choice to $\mathcal{A}$. However, every project that $\mathcal{A}$ initiates reduces the account balance to $x - 1$ (whether or not the latter is positive). Good projects being scarce, there are limits to how many projects the principal can credibly promise. When the balance is at the cap, the account can grow no further; accordingly, the agent takes a project immediately, yielding a balance of $\bar{x} - 1$.

Controlled Budget $(x \leq 0)$

The controlled budget regime is tailored to provide low revenue, low enough to be feasibly provided in aligned equilibrium. When $x < 0$, the agent is over budget, and the principal punishes the agent—more severely the further over budget the agent is— with a freeze, restoring the balance to zero. The continuation contract when the balance is $x = 0$ is our Aligned Optimal Budget.

Definition. *The **Dynamic Capital Budgeting (DCB)** contract $\sigma^{x,\bar{x}}$ is as follows:*

1. *The **Capped Budget** regime:* $x > 0$.

   - *While* $x \in (0, \bar{x})$*:* $\mathcal{P}$ *delegates, and* $\mathcal{A}$ *takes any available good projects and no bad ones. If* $\mathcal{A}$ *initiates a project, the balance jumps from* x *to* $x - 1$*; if* $\mathcal{A}$ *doesn't take a project,* x *drifts according to* $\dot{x} = rx = x > 0$.
   - *When* x *hits* $\bar{x}$*:* $\mathcal{P}$ *delegates, and* $\mathcal{A}$ *takes a project immediately. If* $\mathcal{A}$ *picks up a project, the balance jumps from* $\bar{x}$ *to* $\bar{x} - 1$*; if* $\mathcal{A}$ *doesn't take a project, the balance remains at* $\bar{x}$.

2. *The **Controlled Budget** regime:* $x \leq 0$.

   - *If* $x \in [-1, 0]$*:* $\mathcal{P}$ *freezes for duration* $\log \frac{\omega}{\omega - \underline{\theta}|x|}$*. The Aligned Optimal Budget* $\sigma^{\bar{\tau}}$ *is played thereafter.*

At the physics department's inception, the university allocates a budget of three hires, with a cap of ten. Over time, the physics department searches for candidates. Every time the department finds an appropriate candidate, it hires—and the provost rubber stamps it—spending from the agreed-upon budget. Figure 2 represents one possible realized path of the account balance over time.

The department finds two suitable candidates to hire in its first year; some interest having accrued, the department budget is now at two hires. After the first year, the department enters a dry spell: finding no suitable candidate for six years, the department hires no one. Due to accrued interest, the account budget has increased dramatically from two to over eight hires; furthermore, the increase is exponential. In its ninth year, the account hits the cap and can grow no further. The department can immediately hire up to nine physicists and continue to search (with its remaining budget) or it can hire ten candidates and enact a regime change by the provost. The department chooses to hire one physics professor (irrespective of quality) immediately, and continue to search with a balance of nine.

Over the next few years, the department is flush and hires many professors. First, for three years, the department hits its cap several times, hiring many mediocre candidates. After its eleventh year, the department faces a lucky streak, finding many great physicists over the following years, bringing the budget to one hire. In the next twelve years, the department finds few candidates worth hiring. However, the interest accrual is so slow that the physics department still depletes its budget, in the twenty-eighth year. Throughout this initial phase, the department hires a total of twenty-four physics professors (much more than the account cap of ten).

At this point, the relationship changes permanently. After a temporary hiring freeze, the provost allows the department to resume its search, but follows any hire with a two-year hiring freeze. The relationship is now of a much more conservative character.

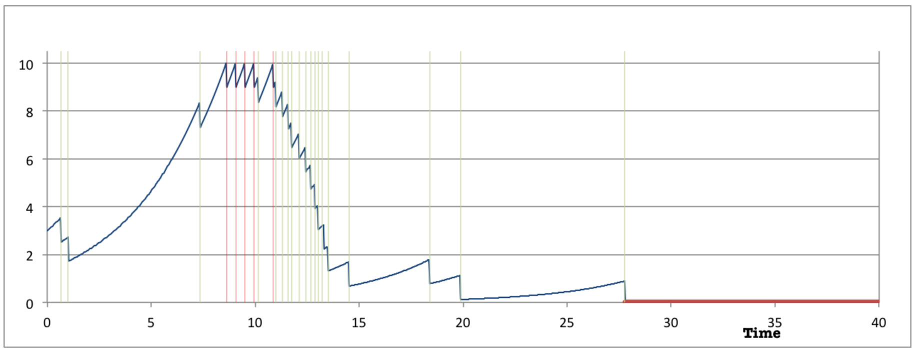

Figure 2: One realization of the balance's path under Controlled Budgeting (with $\bar{x} = 10$). Bad projects are clustered, and the account eventually runs dry.

Notice that bad projects are clustered: immediately after a bad project, the high balance of $\bar{x}-1$ means that the next project is likely bad. Given exponential growth, this effect is stronger the higher is the cap. In the Capped Budget regime, for a given account cap, the balance has non-monotonic profit implications. If the account runs low, there is an increased risk of imminently moving towards the low-revenue Controlled Budget regime. If the account runs high, the principal faces more bad projects in the near future. Observe that Controlled Budgeting is absorbing: once the balance falls low enough—which it eventually does—the agent will never take a bad project again.

Proposition 2. *Fixing an account cap and initial balance* $\bar{x} > x > 0$, *consider the Dynamic Capital Budget contract* $\sigma^{x,\bar{x}}$.

1. $\sigma^{x,\bar{x}}$ *is an equilibrium if and only if it exhibits nonnegative profit at the cap—that is,*

$$\bar{\pi}(\bar{x}) := \pi\Big(\omega + \underline{\theta}\bar{x}, b(\bar{x})\Big) \geq 0.$$

2. *Expected discounted revenue is* $v(x) = \omega + \underline{\theta}x$.

3. *Expected discounted number of bad projects is* $b(x) = b^{\bar{x}}(x)$, *uniquely determined by the delay differential equation*

$$(1+\eta)b(x) = \eta b(x-1) + xb'(x),$$

with boundary conditions:

$$\begin{aligned} b|_{(-\infty,0]} &= 0 \\ b(\bar{x}) - b(\bar{x}-1) &= 1. \end{aligned}$$

Proof. The second point follows from substituting into the v promise-keeping constraint, and noting that (by work in Section 3) $\sigma^{0,\bar{x}}$ yields revenue ω.

The third point follows from our work in Section 3.

For the first part, $v(x) - v(x-1) = [\omega + \underline{\theta}x] - [\omega + \underline{\theta}(x-1)] = \underline{\theta}$ at every x, so that the agent is always indifferent between taking or leaving a bad project. Thus, $\sigma^{x,\bar{x}}$ is an equilibrium if and only if it satisfies principal participation after every history. Revenue is linear, and b is (by work in Section 3) convex. Therefore, profit is concave in x. So, profit is nonnegative for all on-path balances if and only if it is nonnegative at the top. □

Optimality

To gain some intuition as to why the above equilibrium should be optimal, consider how the principal might like to provide different levels of revenue. The case of revenue $v \leq \omega$ is simple: we know that the principal can provide said revenue efficiently—via aligned equilibrium. We also know that other contracts—for instance, Expiring Budget contracts—may yield higher revenue; the key issue is how to provide such higher revenue levels optimally. The principal can provide revenue and incentives via two instruments: (i) punishing the agent for spending, and (ii) rewarding the agent for fiscal restraint. Reminiscent of Ray (2002), the DCB contract backloads costly rewards as much as possible. Subject to satisfying the agent's incentive constraint, the DCB contract uses the minimal punishment possible—i.e. $v - \tilde{v} = \underline{\theta}$—whenever delegating to the agent. Increasing the punishment would accomplish two things, both of them profit-hindering:

1. Following good luck, it would bring the players closer to a low-revenue continuation, where the principal would then have to inefficiently freeze.

2. In accordance with promise-keeping, the increased punishment would then require an accompanying reward for waiting. Following bad luck, this would bring the players to a very high-revenue continuation more quickly, which would necessarily entail more bad projects.

Our main theorem characterizes a profit-maximizing equilibrium of this game. In doing so, we, in fact, achieve a characterization of the whole equilibrium payoff set.

Theorem 2 (DCB Optimality)**.**

1. *There is a cap* $\bar{x}^* \geq 0$ *such that every vector on the Pareto frontier of* $\mathcal{E}^*$ *can be provided by a DCB contract with cap* $\bar{x}^*$*.*

2. *There is a unique initial balance* x^* *such that the DCB contract of cap* $\bar{x}^*$ *and initial balance* x^* *maximizes the principal's value among all equilibria. Moreover, if* $\bar{x}^* > 0$*, then* $x^* > 0$*, and the principal's profit is zero at the top.*

3. *A vector* $(v, b) \geq 0$ *is an equilibrium payoff vector if and only if it yields nonnegative profit to the principal and is (weakly) Pareto-inferior to some DCB contract of cap* $\bar{x}^*$*.*

Proof. Let $B : [0, \bar{v}] \to \mathbb{R}_+$ describe the efficient frontier of the equilibrium value set—i.e. $B(v) = \min\{b : \ (b, v) \in \mathcal{E}^*\}$, where $\bar{v} \geq \omega$ denotes the highest agent value attainable in equilibrium.

In the appendix, we show that

$$B(v) = \begin{cases} 0 = vB'(v) & \text{if } v \in [0, \omega] \\ \dfrac{\eta B(v - \underline{\theta}) + (v - \omega)B'(v)}{1 + \eta} & \text{if } v \in (\omega, \bar{v}) \\ 1 + B(v - \underline{\theta}) & \text{if } v = \bar{v} > \omega, \end{cases}$$

for any $v \in [0, \bar{v}]$. We further show that, if $\bar{v} > \omega$, then $\pi(\bar{v}, B(\bar{v})) = 0$.

We start by focusing on the first claim. By Theorem 1, we know that $\bar{v} \geq \omega$. If all equilibria in the efficient frontier are aligned, $\bar{v} = \omega$, any equilibrium payoff can be supported with initial shutdown followed by the Aligned Optimal Budget. The latter is exactly the DCB contract with $\bar{x}^* = 0$.

If there is a non-aligned efficient equilibrium, then there is some equilibrium that begins with delegation and an immediate bad project–i.e. a gift-giving equilibrium[16]; let $\tilde{v}$ denote the continuation revenue after this bad project. If $\tilde{v} > \omega$, then $\bar{v} > \omega$. If $\tilde{v} \leq \omega$, then principal participation implies that $(1 - \frac{c}{\theta})\omega \geq (c - \underline{\theta})$. In this case, the DCB contract with $\bar{x} = x = 1$ is an equilibrium, so that $\bar{v} > \omega$. Our work with the frontier then implies that $\pi(\bar{v}, B(\bar{v})) = 0$.

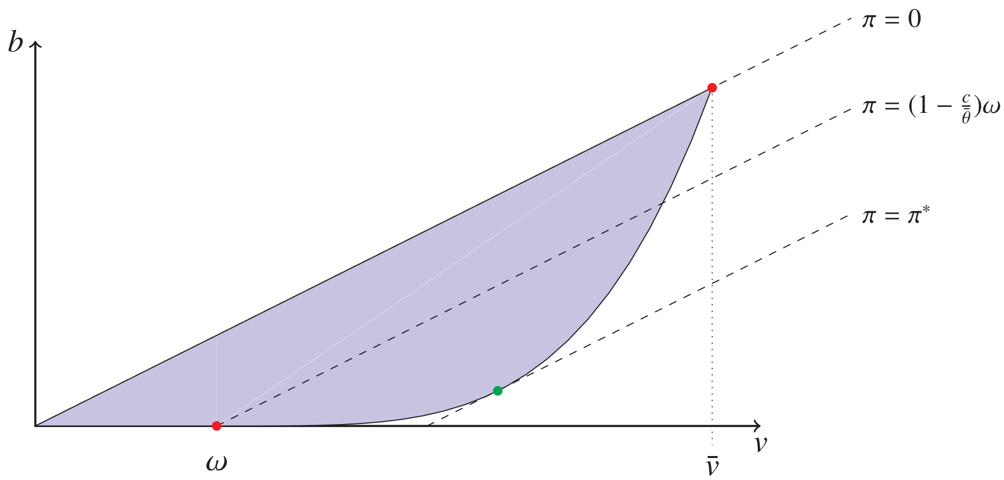

Figure 3: The solid line traces out the frontier B. The equilibrium value set is the convex region between B and the dashed zero-profit line. The dashed lines trace different isoprofits. The green dot highlights the uniquely principal-optimal vector.

Finally, to finish the proof of the first point, all that is needed is to relate the values of $B|_{[\omega,\bar{v}]}$ to a family of DCB contracts (with the same cap). Letting $\bar{x}^* = \frac{\bar{v}-\omega}{\underline{\theta}}$, we now appeal to Proposition 2. For every $x \in [0, \bar{x}^*]$, the DCB contract with a cap of $\bar{x}^*$ and an initial balance of x is an equilibrium providing $(v, b) = (\omega + \underline{\theta}x, B(\omega + \underline{\theta}x))$.

[16]By Lemma 5, in the appendix, the agent never privately mixes in efficient equilibrium.

Toward the second point, first notice that the principal's objective is linear over $\mathcal{E}^*$, so that optimal profit is attained on the graph of $B|_{[\omega,\bar{v}]}$. The result follows trivially if $\bar{x}^* = 0$, there being a unique balance. We now turn our attention to the case of $\bar{x}^* > 0$. Since B is convex, the first-order approach suffices. By the work in Section 3, B is C^1, so that the FOC holds exactly at the optimum, which can be true only for $v > \omega$. Again, by the work in Section 3, B is strictly convex on $[\omega, \bar{v}]$, so that the optimum is unique. Taking an affine transformation, there is a unique optimal balance x^*, which is strictly positive.

Toward the third point, the 'only if' direction follows from the first point and from the principal's participation constraint. For the 'if' direction, take any $(v, b) \in \mathcal{E}^*$. If $v = 0$, then $b = 0$ too (by Principal participation), so that the stage game equilibrium works. If $\bar{v} = \omega$, notice that $(\omega, b_\omega) \in \mathcal{E}^*$[17] where b_ω is (uniquely) such that $\pi(\omega, b_\omega) = 0$. It follows that in this case, the equilibrium value set is all of

$$\mathcal{E}^* = \{(v, b) \geq 0 : \ \pi(v, b) \geq 0, \ v \leq \omega\} = co\{(0,0), \ (\omega, 0), \ (\omega, b_\omega)\}.$$

If $\bar{v} > \omega$, then zero profit at the top implies that

$$\mathcal{E}^* = \{(v, b) \geq 0 : \ \pi(v, b) \geq 0, \ b \geq B(v)\} = co\{(0,0) \cup \text{Graph}[B|_{[\omega,\bar{v}]}]\}.$$

□

The heart of the proof is the characterization of the equilibrium frontier B, formally carried out in the appendix. The overall structure of the argument proceeds as follows. First, allowing for public randomization guarantees us convexity of the equilibrium set frontier. As a consequence, whenever incentivizing picky project adoption by the agent, the principal optimally inflicts the minimum punishment possible. Next, because the principal has no private action or information, we show that the frontier is self-generating and private mixing unnecessary. Collectively, this yields a Bellman equation. Next, we show that initial freeze is inefficient for values

[17] Consider the following profile: the principal plays her continuation strategy exactly as in the Aligned Optimal Budget; whenever the Principal is delegating, the agent takes every good project and takes bad projects with Poisson rate λ such that $\frac{b_\omega}{\omega} = \frac{\lambda}{\eta\bar{\theta}+\lambda\underline{\theta}}$. One readily verifies that this equilibrium strategy profile gives payoffs (ω, b_ω).

above ω, and that bad project adoption is wasteful, except when used to provide value $\bar{v}$.

A direct consequence of our main theorem is that contracts that involve bad projects are necessary for optimality. In what follows, let a *gift-giving equilibrium* be any equilibrium that begins with a certain project: the principal delegates, and the agent initiates a project immediately—no questions asked.

Corollary 2. *The following are equivalent:*

1. *There exists a gift-giving equilibrium.*
2. *Some DCB contract with strictly positive initial balance is an equilibrium.*
3. *Some (non-aligned) DCB contract with strictly positive initial balance strictly outperforms the Aligned Optimal Budget.*

Proof. To see that (1) implies (2), suppose that there exists a gift-giving equilibrium, with an initial gift (i.e. delegation and an immediate project) followed by σ. By the first part of Theorem 2, there is a DCB equilibrium σ'—say, of initial balance x—which (weakly) Pareto dominates σ. If $x > 0$, then (2) follows immediately; if $x = 0$, then σ yields profit $\leq (1 - \frac{c}{\theta})\omega$, and one easily verifies that any DCB contract of cap 1 is an equilibrium. The second part of Theorem 2 ensures that (2) implies (3). Finally, that (3) implies (1) follows from considering a subgame when the account has reached its cap. $\square$

One might suspect that the choice of whether or not to employ bad projects to incentivize picky project adoption by the agent amounts to evaluating a profit tradeoff by the principal. Surprisingly, the principal faces no such tradeoff. Whenever a promise of future bad projects can be credible, it is a necessary component of an optimal contract. Intuition for this result comes from considering DCB contracts with a fixed cap and extremely low positive balances. As the equilibrium frontier is continuously differentiable, increasing an agent's initial balance from zero to a very low $x > 0$ provides a first-order revenue increase and a second-order increase in expected discounted bad projects.

It may appear surprising that one scalar equation can describe the full equilibrium set , even implicitly. This simple structure derives, however, from the paucity of instruments at the principal's disposal.

Note, also, that the DCB contract isn't just optimal; it is essentially uniquely optimal. While there is some flexibility in providing agent values below ω,[18] the optimal way to provide revenue $v \in (\omega, \bar{v})$ is unique. Given that B is strictly convex at v, $B(v) < 1 + B(v - \underline{\theta})$, and $B(v) < vB'(v)$, optimality demands initial delegation paired with picky project adoption and minimum punishment per project. Every principal-optimal contract, therefore, consists of two regimes, the first of which is Capped Budgeting (with the same cap and initial balance). In this sense, dynamic budgeting is not just a useful tool for repeated delegation but, in fact, a necessary one.

Existence of Gift-Giving Equilibria

Theorem 2 provides a complete characterization of the equilibrium payoff set, taking the existence or non-existence of gift-giving equilibria for granted. For any fixed parameters $\eta, \bar{\theta}, \underline{\theta},$ and c, we can determine the existence or non-existence computationally,[19] but we can gain some insight through various sufficient conditions.

Consider the ratio of the net cost of a bad project to the principal's first-best profit,

$$\rho = \frac{c - \underline{\theta}}{\eta(\bar{\theta} - c)}.$$

This ratio is a crude measure of the tradeoff between a bad project today and the principal's future prospects. Even so, for some parameter values, ρ fully resolves the existence question.

If $\rho > 1$, the principal would rather sever the relationship, irrespective of its future value, than admit a bad project: gift-giving is not credible. If $\rho < \frac{\bar{\theta}-\underline{\theta}}{\bar{\theta}}$,

[18] We saw such an example in Section 3, in the "A Smaller Stick" subsection.

[19] Given parameters, we establish an upper bound for the set of possible account caps, $\bar{\bar{x}} = \frac{\omega}{c-\underline{\theta}}$ from principal participation. We numerically solve the delay differential equation in Section 3 of the appendix up to said upper bound. We then explicitly compute the profit at the cap, for each hypothetical cap below the upper bound. In light of Corollary 2, existence of a gift-giving equilibrium is then equivalent to one of these profits being nonnegative.

the profit from the aligned optimal budget more than offsets the net cost of a bad project. Thus, a DCB contract with a cap of 1 is an equilibrium.

While the above sufficient conditions are inconclusive, we note that they reduce the existence problem to the value of ρ whenever the value of a good project dwarfs that of a bad one—i.e. when $\frac{\bar{\theta}}{\underline{\theta}} \approx \infty$. In this case, the inconclusive range $[\frac{\bar{\theta}-\underline{\theta}}{\bar{\theta}}, 1]$ vanishes. In other circumstances, these conditions still yield interpretable predictions. When $c \approx \underline{\theta}$, the principal loses very little from admitting a bad project; thus, ($\rho \approx 0$ tells us) a gift-giving equilibrium exists. When $c \approx \bar{\theta}$, the principal has very little to gain from the relationship's continuation; thus, ($\rho \approx \infty$ tells us) gift-giving cannot happen in equilibrium. Finally, if $\eta \approx \infty$ players are effectively more patient, increasing the future value of their relationship; thus, ($\rho \approx 0$ tells us) gift-giving can be credibly sustained.

Comparative Statics

In light of Theorem 2, a principal-optimal contract is characterized by two elements: how much freedom the principal can credibly give the agent (the cap), and how much freedom the principal chooses to initially give the agent (the initial balance). The first describes the equilibrium payoff set, while the second selects the principal-optimal contract therein. In this subsection, we ask how these features change as the environment the players face varies. As parameters of the model change, and the pool of projects becomes more valuable, the agent enjoys greater sovereignty, with both the balance cap and the optimal initial balance increasing.

Proposition 3. *For any profile of parameters that satisfy Assumptions 1 and 2, define the account cap $\bar{X}(\bar{\theta}, \underline{\theta}, c)$, and the optimal initial account balance $X^*(\bar{\theta}, \underline{\theta}, c)$, as delivered by Theorem 2. Both functions are increasing in the revenue parameters $\bar{\theta}, \underline{\theta}$, and decreasing in the cost parameter c. Moreover, these comparisons are strict in the range where the cap is strictly positive.*

We provide a proof in Section 4 of the appendix. Consider the interesting case, in which the cap is strictly positive. We first analyze a slight increase in a revenue parameter and observe that the profit at the original cap increases. This implies that a slight cap increase maintains the principal's credibility, and it is now an equilibrium. This delivers the first half of our comparative statics result: the new equilibrium has a higher account cap, offering greater flexibility to the agent. As the account cap increases, the frontier of the equilibrium set gets flatter at each account balance. Moreover, as revenue parameters increase, the principal's isoprofit curves can only get steeper. Accordingly, the unique tangency between the equilibrium frontier and an isoprofit occurs at a higher balance. This delivers the remaining comparative statics result: the agent is also given more initial leeway. A similar analysis applies to a cost reduction.

Of particular interest is the role of $\underline{\theta}$ in determining the optimal DCB account structure. On the one hand, the principal suffers less from a bad project when $\underline{\theta}$ is higher; on the other, the agent is more tempted. We show that the former effect always dominates in determining how much freedom the principal optimally gives the agent.

5 Extensions

In this section, we briefly describe some extensions to our model. For simplicity, we restrict attention to the case where gift-giving equilibria exist. The proofs are straightforward and omitted.

Monetary Transfers

We maintain the assumption of limited liability: $\mathcal{A}$ cannot give $\mathcal{P}$ money. If $\mathcal{P}$ can reward $\mathcal{A}$'s fiscal restraint through direct transfers, one of two things happens: (i) nothing changes and money is not used; or (ii) money simply replaces bad projects as a reward if money is more cost-effective. Which is more efficient depends on the relative size of the marginal cost of allowing the agent to initiate bad projects[20] $(c-\underline{\theta})B'(\bar{v})$ and the marginal reward of doing so $\underline{\theta}$. If providing monetary incentives is optimal, a modified DCB contract is used. The cap is raised,[21] and the agent is

[20]This calculation is done using the B from our original model, as characterized in Theorem 2. The condition is correct if $\bar{v} > \omega$; otherwise, it is optimal to use monetary transfers.

[21]The cap is raised to ensure zero profit with the new, more efficient incentivizing technology.

paid a flow of cash whenever his balance is at the cap. This modified DCB contract is reminiscent of the optimal contract in Biais et al. (2010).

Permanent Termination

In many applications, being in a given relationship automatically entails delegating. If a client hires a lawyer, she delegates the choice of hours to be worked. To stop delegating is to terminate the relationship, giving both players zero continuation values.[22] That is, at any moment, the principal must choose between fully delegating and ending the game forever. At first glance, this constraint may seem an additional burden on the principal. However, given our optimal contract (with all freeze backloaded to the Controlled Budget regime), we see that it changes nothing. Indeed, replacing a temporary freeze with stochastic termination[23] leaves payoffs and incentives unchanged.

Agent Replacement

We propose an extension in which the principal has a means to punish the agent without punishing herself: the principal can fire the agent and immediately hire a new agent. The credibility of the threat of replacement takes us far from our leading examples: for instance, the state government cannot sever its relationship with one of its counties.

A fired agent gets a continuation payoff of zero.[24] Every time the principal hires a new agent, she proposes a new contract. We argue below that any inefficiency that the principal faces, as well as any interesting relationship dynamics in the contracts, vanish: in any equilibrium, the principal always delegates to the current agent, who, in turn, exercises fiscal restraint.

[22] The agent could have a positive outside option. As long as it is below $\omega - \underline{\theta}$, the same argument holds.

[23] Keep Capped Budgeting exactly the same. In Controlled Budgeting, replace the duration $\log \frac{\omega}{\omega - \underline{\theta}|x|}$ freeze with a probability $\frac{\underline{\theta}}{\omega}|x|$ termination. As the principal prefers not to terminate the relationship, the randomization must be public.

[24] Again, a positive agent outside option below $\omega - \underline{\theta}$ would change nothing.

In our original model, the only relevant constraint for the principal was the participation constraint. Now, the better outside option for the principal limits the credible promises that she can make to the agent. Being able to terminate this relationship and begin—on her own terms—a new one with a new agent, the principal must expect from this relationship (at any history) at least what she would from a new one—i.e. her optimal value. That $\mathcal{P}$'s outside option and her optimal value coincide forces the principal to claim the same continuation payoff following any history. Finally, note that the principal's first-best profit is attainable: $\mathcal{P}$ delegates, $\mathcal{A}$ initiates only good projects, and every project is followed with agent replacement. In this contract, the principal never freezes: she always delegates to the current agent, who, in turn, adopts only good projects. Although each relationship has, at most, an expected revenue of ω—and, thus, is less profitable than in the optimal DCB contract—the principal's expected total profit across relationships is the first-best $\eta(\bar{\theta} - c)$.

Commitment

If $\mathcal{P}$ has the ability to commit, she can offer $\mathcal{A}$ long-term rewards. In particular, she can offer him tenure (delegation forever) if he exerts fiscal restraint for a long enough time. With full commitment power, slight modifications of our argument show that $\bar{v}$ is the first-best revenue.[25] Guo and Hörner (2014) discuss this case more directly, using the methodology of Spear and Srivastava (1987).

6 Final Remarks

In this paper, we have presented an infinitely repeated instance of the delegation problem. The agent will not represent the principal's interests without being offered dynamic incentives, while the principal cannot credibly commit to long-term rewards.

First, we characterize equilibria that eschew reliance on lenience-based rewards. The principal's hands are tied: she can punish the agent only by limiting her future

[25] We focus on the discrete time setting, so that the agent's first-best outcome is finite.

reliance on his private information, thus harming herself. The Aligned Optimal Budget pairs discerning project adoption with the minimum-length freeze to incentivize it.

Second, we explore the efficiency gains that are possible if bad projects are used as a costly bonus. The promise of future rewards can better incentivize good behavior from the agent, and the value of the future relationship can make such rewards credible for the principal. We characterize the principal-optimal such equilibrium, the Dynamic Capital Budget contract, which comprises two regimes. In the first regime, Capped Budgeting, the agent has an expense account, which grows at the interest rate so long as its balance is below its cap; the principal fully delegates, with every project being financed from the account. The agent takes every available good project; only when at the cap does he adopt projects indiscriminately. Eventually, the account runs dry, and the players transition to the second regime, Controlled Budgeting, wherein they follow the Aligned Optimal Budget. Not only is the DCB contract profit-maximizing, but it in fact traces out the whole equilibrium value set; we note that the analysis and results apply at any fixed discount rate.[26]

The optimal contract suggests rich dynamics for the relationship. Early on, in Capped Budgeting, the relationship is highly productive but low-yield: the agent adopts every good project, but some bad projects as well. The lack of principal commitment limits the magnitude of credible promises, resulting in a transient Capped Budgeting phase. As the relationship matures to Controlled Budgeting, it is high-yield but less productive: the agent adopts only good projects, but some good opportunities go unrealized. In this sense, the relationship drifts toward conservatism.

While our main applications concern organizational economics outside of the firm, we believe that our results also speak to the canonical firm setting.[27] If the relationship between a firm and one of its departments proceeds largely via delegation, then we shed light on the dynamic nature of this relationship. In doing so, we provide a novel foundation for dynamic budgeting within the firm.

[26] In particular, our analysis is not a folk theorem analysis.

[27] The conflict of interest in our model may reflect an empire-building motive on the part of a department, or it may be an expression of the Baumol (1968) sales-maximization principle.

References

Abreu, D., P. Milgrom, and D. Pearce. 1991. "Information and Timing in Repeated Partnerships." *Econometrica* 59 (6):1713–33.

Abreu, D., D. Pearce, and E. Stacchetti. 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58 (5):1041–63.

Alonso, R. and N. Matouschek. 2007. "Relational delegation." *The RAND Journal of Economics* 38 (4):1070–1089.

Ambrus, A. and G. Egorov. 2013. "Delegation and Nonmonetary Incentives." Working paper, Duke University & Northwestern Kellogg.

Armstrong, M. and J. Vickers. 2010. "A Model of Delegated Project Choice." *Econometrica* 78 (1):213–44.

Baumol, W. 1968. "On the Theory of Oligopoly." *Economica* 25 (99):187–98.

Biais, B., T. Mariotti, J. Rochet, and S. Villeneuve. 2010. "Large Risks, Limited Liability, and Dynamic Moral Hazard." *Econometrica* 78 (1):73–118.

Casella, A. 2005. "Storable Votes." *Games and Economic Behavior* 51:391–419.

Chassang, S. 2010. "Building Routines: Learning, Cooperation, and the Dynamics of Incomplete Relational Contracts." *American Economic Review* 100 (1):448–65.

Clementi, G. and H. Hopenhayn. 2006. "A Theory of Financing Constraints and Firm Dynamics." *The Quarterly Journal of Economics* 121 (1):1131–1175.

Compte, Olivier. 1998. "Communication in repeated games with imperfect private monitoring." *Econometrica* :597–626.

Espino, E., J. Kozlowski, and J. M. Sanchez. 2013. "Too big to cheat: efficiency and investment in partnerships." *Federal Reserve Bank of St. Louis Working Paper Series* (2013-001).

Frankel, A. 2011. “Discounted Quotas.” Working paper, Stanford University.

———. 2013. “Delegating Multiple Decisions.” Working paper, Chicago Booth.

———. 2014. “Aligned Delegation.” *American Economic Review* 104 (1):66–83.

Guo, Y. 2014. “Dynamic Delegation of Experimentation.” Working paper, Northwestern University.

Guo, Y. and J. Hörner. 2014. “Dynamic Mechanisms without Money.” Working paper, Northwestern University & Yale University.

Harrington, Joseph E and Andrzej Skrzypacz. 2011. “Private monitoring and communication in cartels: Explaining recent collusive practices.” *The American Economic Review* 101 (6):2425–2449.

Hauser, C. and H. Hopenhayn. 2008. “Trading Favors: Optimal Exchange and Forgiveness.” Working paper, Collegio Carlo Alberto & UCLA.

Holmström, B. 1984. “On the Theory of Delegation.” In *Bayesian Models in Economic Theory*. New York: North-Holland, 115–41.

Jackson, M. and H. Sonnenschein. 2007. “Overcoming Incentive Constraints by Linking Decisions.” *Econometrica* 75 (1):241–57.

Levin, J. 2003. “Relational incentive contracts.” *The American Economic Review* 93 (3):835–857.

Li, J., N. Matouschek, and M. Powell. 2015. “Power Dynamics in Organizations.” Working paper, Northwestern University.

Li, Jin and Niko Matouschek. 2013. “Managing conflicts in relational contracts.” *the American Economic Review* 103 (6):2328–2351.

Malcomson, James M. 2010. “Relational Incentive Contracts.” .

Malenko, A. 2013. “Optimal Dynamic Capital Budgeting.” Working paper, MIT.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic theory*. Oxford: Oxford University Press.

Möbius, M. 2001. “Trading Favors.” Working paper, Harvard University.

Pearce, D. G. and E. Stacchetti. 1998. “The interaction of implicit and explicit contracts in repeated agency.” *Games and Economic Behavior* 23 (1):75–96.

Ray, Debraj. 2002. “The Time Structure of Self-Enforcing Agreements.” *Econometrica* 70 (2):547–582.

Spear, S. and S. Srivastava. 1987. “On Repeated Moral Hazard with Discounting.” *Review of Economic Studies* 54 (4):599–617.

Online Appendix: Not for Publication

This online appendix provides formal proof for the results on which the main text draws. First, we provide a characterization of the equilibrium payoff set of the discrete-time game. Then, we provide auxiliary computations concerning Expiring Budget contract. Next, we prove several useful properties of the Delay Differential Equation which characterizes the frontier of the limit equilibrium set. Finally, we derive comparative statics results for the optimal cap and initial balance of a Dynamic Capital Budget contract.

1 APPENDIX: Characterizing the Equilibrium Value Set

In the current section, we characterize the equilibrium value set in our discrete time repeated game. As in the main text, we find it convenient to study payoffs in terms of *agent value* and *bad projects*. Accordingly, for any strategy profile σ, we let

$$\begin{aligned} v(\sigma) &= \mathbb{E}^{\sigma}\left[\sum_{k=0}^{\infty} \delta^k \mathbf{1}_{\{\text{a project is picked up in period } k\}} \theta_k\right]; \\ b(\sigma) &= \mathbb{E}^{\sigma}\left[\sum_{k=0}^{\infty} \delta^k \mathbf{1}_{\{\text{a project is picked up in period } k\}} 1_{\theta_k=\underline{\theta}}\right]. \end{aligned}$$

Below, we will analyze the public perfect equilibrium (PPE) value set,

$$\mathcal{E}^* = \{(v(\sigma), b(\sigma)) : \ \sigma \text{ is a PPE}\} \subseteq \mathbb{R}^2_+.$$

1.1 Self-Generation

To describe the equilibrium value set $\mathcal{E}$, we rely heavily on the machinery of Abreu, Pearce, and Stacchetti (1990), APS. To provide the players a given value $y = (v, b)$ from today onward, we factorize it into a (possibly random) choice of what happens today, and what the continuation will be starting tomorrow. What happens today depends on the probability (p) that the principal delegates, the probability ($\bar{a}$) of project adoption if a project is good, and the probability ($\underline{a}$) of project adoption if a project is bad. The continuation values may vary based on what happens today: the principal may choose to freeze ($\check{y}$), the principal may delegate and agent may take a project ($\tilde{y}$), or the principal may delegate and agent may not take a project (y'). Since the principal doesn't observe project types, these are the only three public outcomes.

We formalize this factorization in the following definition and theorem.

Definition 2. *Given $Y \subseteq \mathbb{R}^2$:*

- *Say $y \in \mathbb{R}^2$ is* ***purely enforceable*** *w.r.t. Y if there exist $p, \bar{a}, a \in [0,1]$ and $\check{y}, \tilde{y}, y' \in Y$ such that:*[28]

  1. *(Promise keeping):*

$$\begin{aligned} y &= (1-p)\delta\check{y} + ph\left\{\bar{a}\left[(\bar{\theta},0)+\delta\tilde{y}\right]+(1-\bar{a})\delta y'\right\} \\ &\quad +p(1-h)\left\{a\left[(\underline{\theta},1)+\delta\tilde{y}\right]+(1-a)\delta y'\right\} \\ &= (1-p)\delta\check{y} + p\Big\{h\bar{a}\left[(\bar{\theta},0)+\delta(\tilde{y}-y')\right]+(1-h)a\left[(\underline{\theta},1)+\delta(\tilde{y}-y')\right]+\delta y'\Big\}. \end{aligned}$$

  2. *(Incentive-compatibility):*

$$\begin{aligned} p &\in \arg\max_{\hat{p}\in[0,1]} \quad \hat{p}\Big\{h\bar{a}\left[(\bar{\theta}-c)+\delta[\pi(\tilde{y})-\pi(y')]\right]+(1-h)a\left[(\underline{\theta}-c)+\delta[\pi(\tilde{y})-\pi(y')]\right]+\delta\pi(y')-\delta\pi(\\ \bar{a} &\in \arg\max_{\hat{a}\in[0,1]} \quad \hat{a}\left\{\bar{\theta}+\delta[v(\tilde{y})-v(y')]\right\}, \\ a &\in \arg\max_{\hat{a}\in[0,1]} \quad \hat{a}\left\{\underline{\theta}+\delta[v(\tilde{y})-v(y')]\right\}. \end{aligned}$$

- *Say $y \in \mathbb{R}^2$ is* ***enforceable*** *w.r.t. Y if there exists a Borel probability measure μ on $\mathbb{R}^2$ such that*

  1. *$y = \int_{\mathbb{R}^2} \hat{y}\, d\mu(\hat{y})$.*
  2. *$\hat{y}$ is purely enforceable almost surely with respect to $\mu(\hat{y})$.*

- *Let $W(Y) := \{y \in \mathbb{R}^2 : \ y \text{ is enforceable with respect to } Y\}$.*

- *Say $Y \subseteq \mathbb{R}^2$ is* ***self-generating*** *if $Y \subseteq W(Y)$.*

[28]With a slight abuse of notation, for a given $y = (y_1, y_2) \in \mathbb{R}^2$, we will let $v(y) := y_1$.

Adapting methods from Abreu, Pearce, and Stacchetti (1990), one can readily characterize $\mathcal{E}$ via self-generation, through the following collection of results.

Lemma 2. *Let W be as defined above.*

- *The set operator $W : 2^{\mathbb{R}^2} \longrightarrow 2^{\mathbb{R}^2}$ is monotone.*
- *Every bounded, self-generating $Y \subseteq \mathbb{R}^2$ is a subset of $\mathcal{E}^*$.*
- *$\mathcal{E}^*$ is the largest bounded self-generating set.*
- *$W(\mathcal{E}^*) = \mathcal{E}^*$.*
- *Let $Y_0 \subseteq \mathbb{R}^2$ be any bounded set with*[29] *$\mathcal{E}^* \subseteq W(Y_0) \subseteq Y_0$. Define the sequence $(Y_n)_{n=1}^{\infty}$ recursively by $Y_n := W(Y_{n-1})$ for each $n \in \mathbb{N}$. Then $\bigcap_{n=1}^{\infty} Y_n = \mathcal{E}^*$.*

1.2 A Cleaner Characterization

In light of the above, understanding the operator W will enable us to fully describe $\mathcal{E}^*$. That said, the definition of W is somewhat cumbersome. For the remainder of the current section, we work to better understand it.

Before doing anything else, we restrict attention to a useful domain for the map W.

Notation. *Let $\mathcal{Y} := \{Y \subseteq \mathbb{R}^2_+ : \ \vec{0} \in Y, \ \ Y$ is compact and convex, and $\pi|_Y \geq 0\}$.*

We need to work only with potential value sets in $\mathcal{Y}$. Indeed, the feasible set $\bar{\mathcal{E}}$ belongs to $\mathcal{Y}$, and it is straightforward to check that W takes elements of $\mathcal{Y}$ to $\mathcal{Y}$. Since $\mathcal{Y}$ is closed under intersections, we then know from the last bullet of the result above, Lemma 2, that $\mathcal{E} \in \mathcal{Y}$.

In seeking a better description of W, the following auxiliary definitions are useful.

Definition 3. *Given $a \in [0, 1]$:*

- *Say $y \in \mathbb{R}^2_+$ is a-**Pareto enforceable** w.r.t. Y if there exist $\tilde{y}, y' \in Y$ such that:*

  1. *(Promise keeping):*

$$y = h\left[(\bar{\theta}, 0) + \delta(\tilde{y} - y')\right] + (1 - h)a\left[(\underline{\theta}, 1) + \delta(\tilde{y} - y')\right] + \delta y'.$$

[29] This can be ensured, for instance, by letting $\mathcal{E}^*$ contain the feasible set, scaled by $\frac{1}{1-\delta}$.

2. *(Agent incentive-compatibility):*

$$1 \in \arg\max_{\hat{a} \in [0,1]} \quad \hat{a}\left\{\bar{\theta} + \delta[v(\tilde{y}) - v(y')]\right\},$$
$$a \in \arg\max_{\hat{a} \in [0,1]} \quad \hat{a}\left\{\underline{\theta} + \delta[v(\tilde{y}) - v(y')]\right\}.$$

3. *(Principal participation):* $\pi(y) \geq 0$.

- *Let* $W_a(Y) := \{y \in \mathbb{R}^2_+ : \ y$ *is a-Pareto enforceable w.r.t.* $Y\}$.
- *Let* $W_f(Y) := \delta Y$.
- *Let* $\hat{W}(Y) := W_f(Y) \cup \bigcup_{a \in [0,1]} W_a(Y)$. *If* Y *is compact, then so is* $\hat{W}(Y)$.[30]

The set $\hat{W}(Y)$ captures the enforceable (without public randomizations) values w.r.t. Y if:

1. The principal uses a pure strategy.
2. We relax principal IC to a participation constraint.
3. If the principal delegates and the project is good, then the agent takes the project.

The following proposition shows that, for the relevant $Y \in \mathcal{Y}$, it is without loss to focus on $co\hat{W}$ instead of W. The result is intuitive. The first two points are without loss because the principal's choices are observable. Toward (1), her private mixing can be replaced with public mixing with no effect on $\mathcal{A}$'s incentives. Toward (2), if the principal faces nonnegative profits with any pure action, she can be incentivized to take said action with stage Nash (min-max payoffs) continuation following the other choice. Toward (3), the agent's private mixing isn't (given (2)) important for the principal's IC, and so we can replace it with public mixing between efficient (i.e. no good project being passed up) first-stage play and an initial freeze.

Lemma 3. *If* $Y \in \mathcal{Y}$, *then* $W(Y) = co\hat{W}(Y)$.

Proof. First, notice that $\delta Y \subseteq W(Y) \cap co\hat{W}(Y)$. It is a subset of the latter by construction, and of the former by choosing $\tilde{y} = y' = \vec{0}$, $p = 1, \bar{a} = a = 1$, and letting $\check{y}$ range over Y.

[30] Indeed, it is the union of δY and a projection of the compact set $\{(a, y) \in [0,1] \times \mathbb{R}^2 : y$ is a-Pareto enforceable w.r.t. $Y\}$.

Take any $y \in \hat{W}(Y)$ that isn't in δY. So y is a-Pareto enforceable w.r.t. Y for some $a \in [0,1]$, say witnessed by $\tilde{y}, y' \in Y$. Letting $p = 0$, $\bar{a} = 1$, and $\check{y} = \vec{0} \in Y$, it is immediate that $p, \bar{a}, a \in [0,1]$ and $\check{y}, \tilde{y}, y' \in Y$ witness y being purely enforceable w.r.t. Y. Therefore, $y \in W(Y)$. So $\hat{W}(Y) \subseteq W(Y)$. The latter being convex, $co\hat{W}(Y) \subseteq W(Y)$ as well.

Take any extreme point y of $W(Y)$ which isn't in δY. Then y must be purely enforceable w.r.t. Y, say witnessed by $p, \bar{a}, a \in [0,1]$ and $\check{y}, \tilde{y}, y' \in Y$. First, if $p\bar{a} = 0$, then[31]

$$y = (1-p)\delta\check{y} \in co\{\vec{0}, \delta\check{y}\} \subseteq \delta Y \subseteq \hat{W}(Y).$$

Now suppose $p\bar{a} > 0$, and define[32] $a^p := \dfrac{a}{\bar{a}}$ and

$$y^p := h\left[(\bar{\theta}, 0) + \delta(\tilde{y} - y')\right] + (1-h)a^p\left[(\underline{\theta}, 1) + \delta(\tilde{y} - y')\right] + \delta y'.$$

Observe that $\tilde{y}, y'$ witness $y^p \in W_{a^p}(Y)$:

1. Promise keeping follows immediately from the definition of y^p.

2. Agent IC follows from agent IC in enforcement of y, and from the fact that incentive constraints are linear in action choices. As $\bar{a} > 0$ was optimal, $\bar{a}^p = 1$ is optimal here as well.

3. Principal participation follows from principal IC in enforcement of y, and from the fact that $\pi(\check{y}) \geq 0$ because $\pi|_Y \geq 0$.

Therefore $y^p \in W_{a^p}(Y)$, from which it follows that $y = (1-p)\delta\check{y} + p\bar{a}y^p \in co\{\delta\check{y}, y^p, \vec{0}\} \subseteq \hat{W}(Y)$.

As every extreme point of $W(Y)$ belongs to $\hat{W}(Y)$, all of $W(Y)$ belongs to the closed convex hull of $\hat{W}(Y)$, which is just $co\hat{W}(Y)$.[33] □

In view of the above proposition, we now only have to consider the much simpler map $\hat{W}$. As the following lemma shows, we can even further simplify, by showing that there is never a need to offer excessive punishment. That is, it is without loss to (1) make the agent's IC constraint (to resist bad projects) bind if he is being discerning, and (2) not respond to the agent's choice if he is being indiscriminate.

[31] If $p\bar{a} = 0$, then either $p = 0$ or $\bar{a} = 0$. If $\bar{a} = 0$, then agent IC implies $a = 0$. So either $p = 0$ or $a = \bar{a} = 0$; in either case, promise keeping then implies $y = (1-p)\delta\check{y}$.

[32] Since $a \leq \bar{a}$ by IC, we know $a^p \in [0,1]$.

[33] The disappearance of the qualifier "closed" comes from Carathéodory's theorem, since $\hat{W}(Y)$ is compact in Euclidean space.

Lemma 4. *Fix $a \in [0,1]$, $Y \in \mathcal{Y}$, and $y \in \mathbb{R}^2$:*

Suppose $a < 1$. Then $y \in W_a(Y)$ if and only if there exist $\tilde{z}, z' \in Y$ such that:

1. *(Promise keeping):*

$$y = h\left[(\bar{\theta}, 0) + \delta(\tilde{z} - z')\right] + (1-h)a\left[(\underline{\theta}, 1) + \delta(\tilde{z} - z')\right] + \delta z'.$$

2. *(Agent* ***exact*** *incentive-compatibility):*

$$\delta[v(z') - v(\tilde{z})] = \underline{\theta}.$$

3. *(Principal participation): $\pi(y) \geq 0$.*

Suppose $a = 1$. Then $y \in W_a(Y)$ if and only if there exists $z' \in Y$ such that:

1. *(Promise keeping):*

$$y = h(\bar{\theta}, 0) + (1-h)(\underline{\theta}, 1) + \delta z'$$

2. *(Principal participation): $\pi(y) \geq 0$.*

Proof. In the first case, the "if" direction is immediate from the definition of W_a. In the second, it is immediate once we apply the definition of W_1 with $\tilde{z} = z'$. Now we proceed to the "only if" direction.

Consider any $y \in W_a(Y)$, with $\tilde{y}, y'$ witnessing a-Pareto enforceability. Define

$$\bar{y} := [h + a(1-h)]y' + (1-h)(1-a)\tilde{y} \in Y.$$

So $\bar{y}$ is the on-path expected continuation value.

In the case of $a < 1$, define

$$\begin{aligned} q &:= \frac{\underline{\theta}}{\delta[v(y') - v(\tilde{y})]} \quad (\in [0,1], \text{ by IC}) \\ \tilde{z} &:= (1-q)\bar{y} + q\tilde{y} \\ z' &:= (1-q)\bar{y} + qy'. \end{aligned}$$

By construction, $\delta[v(z') - v(\tilde{z})] = \underline{\theta}$, as desired.[34]

[34] In the case of $a \in (0,1)$, $q = 1$ (by agent IC), so that $\tilde{z} = \tilde{y}$ and $z' = y'$. The real work was needed for the case of $a = 0$.

In the case of $a = 1$, let $z' := \bar{y}$ and $\tilde{z} := z'$.

Notice that $\tilde{z}, z'$ witness $y \in W_a(Y)$. Promise keeping comes from the definition of $\bar{y}$, principal participation comes from the hypothesis that $W_a(Y) \ni y$, and IC (exact in the case of $a < 1$) comes by construction. □

In the first part of the lemma, $\delta[v(y') - v(\tilde{y})] \in [\underline{\theta}, \bar{\theta}]$ has been replaced with $\delta[v(y') - v(\tilde{y})] = \underline{\theta}$. That is, it is without loss to make the agent's relevant incentive constraint—to avoid taking bad projects—bind. This follows from the fact that $Y \supseteq co\{\tilde{y}, y'\}$. The second part of the lemma says that, if the agent isn't being at all discerning, nothing is gained from disciplining him.

The above lemma has a clear interpretation, familiar from the Aligned Optimal Budget: without loss of generality, the principal uses the minimal possible punishment. The lemma also yields the following:

Lemma 5. *Suppose $a \in (0, 1)$, $Y \in \mathcal{Y}$, and $y \in W_a(Y)$. Then there is some $y^* \in W_0$ such that*

$$v(y^*) = v(y) \text{ and } b(y^*) < b(y).$$

That is, $y_1^* = y_1$ and $y_2^* < y_2$.

Proof. Appealing to Lemma 4, there exist $\tilde{z}, z' \in Y$ such that:

1. (Promise keeping):

$$y = h\left[(\bar{\theta}, 0) + \delta(\tilde{z} - z')\right] + (1 - h)a\left[(\underline{\theta}, 1) + \delta(\tilde{z} - z')\right] + \delta z'$$

2. (Agent **exact** incentive-compatibility):

$$\delta[v(z') - v(\tilde{z})] = \underline{\theta}$$

3. (Principal participation): $\pi(y) \geq 0$.

Given agent exact IC, we know $v(z') > v(\tilde{z})$. Let $\tilde{z}^* := \left(v(\tilde{z}), \min\left\{b(\tilde{z}), \frac{v(\tilde{z})}{v(z')} b(z')\right\}\right)$. As either $\tilde{z}^* = \tilde{z}$ or $\tilde{z}^* \in co\{\vec{0}, z'\}$, we have $z^* \in Y$.

Let $y^* := h\left[(\bar{\theta}, 0) + \delta(\tilde{z}^* - z')\right] + \delta z'$. Then

$$\begin{aligned} v(y) - v(y^*) &= (1-h)a\left\{\underline{\theta} + \delta[v(\tilde{z}^*) - v(z')]\right\} - h\delta[v(\tilde{z}^*) - v(\tilde{z})] \\ &= (1-h)a\left\{\underline{\theta} + \delta[v(\tilde{z}) - v(z')]\right\} - h\delta 0 \\ &= 0, \end{aligned}$$

while

$$\begin{aligned} b(y) - b(y^*) &= (1-h)a\{1 + \delta[b(\tilde{z}) - b(z')]\} - h\delta[b(\tilde{z}^*) - b(\tilde{z})] \\ &= (1-h)a\{1 + \delta[b(\tilde{z}^*) - b(z')] + \delta[b(\tilde{z}) - b(\tilde{z}^*)]\} - h\delta[b(\tilde{z}^*) - b(\tilde{z})] \\ &= (1-h)a\{1 + \delta[b(\tilde{z}^*) - b(z')]\} + [h + (1-h)a]\delta[b(\tilde{z}) - b(\tilde{z}^*)] \\ &\geq (1-h)a \\ &> 0. \end{aligned}$$

Now, notice that $\tilde{z}^*, z'$ witness $y^* \in W_0(Y)$. Promise keeping holds by fiat, agent IC holds because $v(\tilde{z}^*) = v(\tilde{z})$ by construction, and principal participation follows from

$$\pi(y^*) - \pi(y) = -\pi(0, b(y) - b(y^*)) > 0.$$

□

The above lemma is a strong bang-bang result. It isn't just sufficient to restrict attention to equilibria with no private mixing; it is necessary too. Any equilibrium in which the agent mixes on-path is Pareto dominated.

1.3 Self-Generation for Frontiers

Through Lemmata 3, 4, and 5, we greatly simplified analysis of the APS operator W applied to the relevant value sets. In the current section, we profit from that simplification in characterizing the efficient frontier of $\mathcal{E}^*$. Before we can do that, however, we have to make a small investment in some new definitions. We then translate the key results of the previous subsection into results about the efficient frontier of the equilibrium set.

Notation. *Let $\mathcal{B}$ denote the space of functions $B : \mathbb{R} \longrightarrow \mathbb{R}_+ \cup \{\infty\}$ such that: (1) B is convex, (2) $B(0) = 0$, (3) B's proper domain $dom(B) := B^{-1}(\mathbb{R})$ is a compact subset of $\mathbb{R}_+$, (4) B is continuous on $dom(B)$, and (5) $\pi(v, B(v)) \geq 0$ for every $v \in dom(B)$.*

Just as $\mathcal{Y}$ is the relevant space of value sets, $\mathcal{B}$ is the relevant space of frontiers of value sets.

Notation. *For each $Y \in \mathcal{Y}$, define the **efficient frontier function** of Y:*

$$\begin{aligned} B_Y : \mathbb{R} &\longrightarrow \mathbb{R}_+ \cup \{\infty\} \\ v &\longmapsto \min\{b \in \mathbb{R}_+ : \quad (v, b) \in Y\} \end{aligned}$$

It is immediate that for $Y \in \mathcal{Y}$, the function B_Y belongs to $\mathcal{B}$.

Notation. *Define the following functions:*

$$\begin{aligned} T : \mathcal{B} &\longrightarrow \mathcal{B} \\ \hat{B} &\longmapsto B_{W\left(co[graph(\hat{B})]\right)} = B_{co\hat{W}\left(co[graph(\hat{B})]\right)}, \\ T_f : \mathcal{B} &\longrightarrow \mathcal{B} \\ \hat{B} &\longmapsto B_{W_f\left(co[graph(\hat{B})]\right)} = B_{\delta\left(co[graph(\hat{B})]\right)}, \\ \text{and for } 0 \leq a \leq 1, \quad T_a : \mathcal{B} &\longrightarrow \mathcal{B} \\ \hat{B} &\longmapsto B_{W_a\left(co[graph(\hat{B})]\right)}. \end{aligned}$$

These objects are not new. The map T [resp. T_f, T_a] is just a repackaging of W [resp. W_f, W_a], made to operate on frontiers of value sets, rather than on value sets themselves.

As it turns out, we really only need Lemmata 3, 4, and 5 to simplify our analysis of T, which in turn helps us characterize the efficient frontier of $\mathcal{E}^*$. We now proceed along these lines.

The following lemma is immediate from the definition of the map $Y \mapsto B_Y$.

Lemma 6. *If $\{Y_i\}_{i\in\mathbb{I}} \subseteq \mathcal{Y}$, then $B_{\overline{co}[\bigcup_{i\in\mathbb{I}} Y_i]}$ is the convex lower envelope[35] of $\inf_{i\in\mathbb{I}} B_{Y_i}$.*

The following proposition is the heart of our main characterization of the set $\mathcal{E}^*$'s frontier. It amounts to a complete description of the behavior of T.

Proposition 4. *Fix any $B \in \mathcal{B}$ and $v \in \mathbb{R}$. Then:*

1. $TB = cvx\left[\min\left\{T_f B, T_0 B, T_1 B\right\}\right]$.

[35]The **convex lower envelope** of a function $\check{B}$ is cvx$\check{B}$,the largest convex upper-semicontinuous function below it. Equivalently, cvx$\check{B}$ is the pointwise supremum of all affine functions below $\check{B}$.

2. *For* $i \in \{f, 0, 1\}$,

$$T_i B(v) = \begin{cases} \check{T}_i B(v) & \textit{if } \pi\left(v, \check{T}_i^{\Delta} B(v)\right) \geq 0 \\ \infty & \textit{otherwise,} \end{cases}$$

where

$$\begin{aligned} \check{T}_f B(v) &:= \delta B\left(\frac{v}{\delta}\right) \\ \check{T}_0 B(v) &:= \delta\left[hB\left(\frac{v - \theta_E}{\delta}\right) + (1-h)B\left(\frac{v - [\theta_E - \underline{\theta}]}{\delta}\right)\right] \\ \check{T}_1 B(v) &:= (1-h) + \delta B\left(\frac{v - \theta_E}{\delta}\right) \end{aligned}$$

Proof. That $TB = \text{cvx}\left[\min\left\{T_f B,\ \inf_{a \in [0,1]} T_a B\right\}\right]$ is a direct application of Lemma 6. Then, appealing to Lemma 5, $T_a B \geq T_0 B$ for every $a \in (0, 1)$. This proves the first point.

In what follows, let $Y := co[\text{graph}(B)]$ so that $TB = B_{W(Y)}$.

- Consider any $y \in W_0(Y)$:
  Lemma 4 delivers $\tilde{z}, z' \in Y$ such that

  $$\begin{aligned} y &= h\left[(\bar{\theta}, 0) + \delta(\tilde{z} - z')\right] + \delta z', \\ \underline{\theta} &= \delta[v(z') - v(\tilde{z})]. \end{aligned}$$

  Rewriting with $\tilde{z} = (\tilde{v}, \tilde{b})$ and $z' = (v', b')$, and rearranging yields:

  $$\begin{aligned} \underline{\theta} &= \delta[v' - \tilde{v}] \\ (v, b) &= h\left[(\bar{\theta}, 0) + \delta(\tilde{v} - v', \tilde{b} - b')\right] + \delta(v', b') \\ &= h\left(\bar{\theta} - \underline{\theta},\ \ \delta[\tilde{b} - b']\right) + \delta(v', b') \\ &= \left(\theta_E - \underline{\theta} + \delta v',\ \ h\delta\tilde{b} + (1-h)\delta b'\right) \end{aligned}$$

  Solving for the agent values yields

  $$v' = \frac{v - [\theta_E - \underline{\theta}]}{\delta} \text{ and } \tilde{v} = v' - \delta^{-1}\underline{\theta} = \frac{v - \theta_E}{\delta}.$$

So given any $v \in \mathbb{R}_+$:

$$\begin{aligned}
T_0 B(v) \quad &= \quad \inf_{b,\tilde{b},b'} \ b \\
\text{s.t.} \quad & \pi(v,b) \geq 0, \ \ b = \delta\left[h\tilde{b} + (1-h)b'\right], \text{ and } \left(\frac{v - [\theta_E - \underline{\theta}]}{\delta}, \ \ b'\right), \left(\frac{v - \theta_E}{\delta}, \ \ \tilde{b}\right) \in Y \\
&= \quad \inf_{b,\tilde{b},b'} \ b = \delta\left[h\tilde{b} + (1-h)b'\right] \\
\text{s.t.} \quad & \pi(v,b) \geq 0 \text{ and } \left(\frac{v - [\theta_E - \underline{\theta}]}{\delta}, \ \ b'\right), \left(\frac{v - \theta_E}{\delta}, \ \ \tilde{b}\right) \in Y \\
&= \quad \begin{cases} b = \delta\left[hB\left(\frac{v-\theta_E}{\delta}\right) + (1-h)B\left(\frac{v-[\theta_E-\underline{\theta}]}{\delta}\right)\right] & \text{if } \pi(v,b) \geq 0, \\ \infty & \text{otherwise.} \end{cases}
\end{aligned}$$

- Consider any $y \in W_1(Y)$:
  Lemma 4 now delivers $z' = (v', b') \in Y$ such that

$$y = h(\bar{\theta}, 0) + (1-h)(\underline{\theta}, 1) + \delta z',$$

which can be rearranged to

$$(v, b) = (\theta_E + \delta v', \ \ (1-h) + \delta b').$$

So given any $v \in \mathbb{R}_+$:

$$\begin{aligned}
T_1 B(v) \quad &= \quad \inf_{b,b'} b \\
\text{s.t.} \quad & \pi(v,b) \geq 0, \ \ b = (1-h) + \delta b', \text{ and } \left(\frac{v - \theta_E}{\delta}, \ \ b'\right) \in Y \\
&= \quad \inf_{\tilde{b},b'} b = (1-h) + \delta b' \\
\text{s.t.} \quad & \pi(v,b) \geq 0 \text{ and } \left(\frac{v - \theta_E}{\delta}, \ \ b'\right) \in Y \\
&= \quad \begin{cases} b = \delta B\left(\frac{v-\theta_E}{\delta}\right) & \text{if } \pi(v,b) \geq 0, \\ \infty & \text{otherwise.} \end{cases}
\end{aligned}$$

- Lastly, given any $v \in \mathbb{R}_+$:

$$\begin{aligned} T_f B(v) &= \inf_{b,b'} b \\ \text{s.t.} &\quad \pi(v,b) \geq 0,\ \ b = \delta b',\ \text{and}\ \left(\frac{v}{\delta}, b'\right) \in Y \\ &= \begin{cases} b = \delta B\left(\frac{v}{\delta}\right) & \text{if } \pi(v,b) \geq 0, \\ \infty & \text{otherwise.} \end{cases} \end{aligned}$$

□

1.4 The Efficient Frontier

In this subsection, we characterize the frontier $B_{\mathcal{E}^*}$ of the equilibrium value set. We first translate APS's self-generation to the setting of frontiers. This, along with Proposition 4 delivers our Bellman equation, Corollary 3. Then, in Theorem 1, we characterize the discrete time equivalent of the Aligned Optimal Budget. Finally, in Theorem 2, we fully characterize the frontier $B_{\mathcal{E}^*}$.

Theorem 3. *Suppose $Y \in \mathcal{Y}$ with $W(Y) = Y$. Then $TB_Y = B_Y$.*

Proof. First, because W is monotone, $W(co[\text{graph}(B_Y)]) \subseteq W(Y) = Y$. Thus the efficient frontier of the former is higher than that of the latter. That is, $TB_Y \geq B_Y$.

Now take any $v \in \text{dom}(B_Y)$ such that $y := (v, B_Y(v))$ is an extreme point of Y. We want to show that $TB_Y(v) \leq B_Y(v)$.

By Lemma 3, $y \in W_f(Y) \cup \bigcup_{a\in[0,1]} W_a(Y)$.

- If $y \in W_f(Y)$, then $\frac{y}{\delta}, \vec{0} \in Y$, so that the extreme point y must be equal to $\vec{0}$. But in this case, $TB_Y(v) = TB_Y(0) = 0 = B_Y(0) = B_Y(v)$.
- If $y \in W_a(Y)$ for some $a \in [0,1]$, say witnessed by $\tilde{y}, y' \in Y$, then let

$$\begin{aligned} \tilde{z} &:= (v(\tilde{y}), B_Y(v(\tilde{y}))) \\ \tilde{z}' &:= (v(y'), B_Y(v(y'))) \\ z &:= h\left[(\bar{\theta}, 0) + \delta(\tilde{z} - z')\right] + (1-h)a\left[(\underline{\theta}, 1) + \delta(\tilde{z} - z')\right] + \delta z'. \end{aligned}$$

Then

$$\begin{aligned} b(z) &= (1-h)a + [h + (1-h)a]\delta B_Y(v(\tilde{y})) + (1-h)(1-a)\delta B_Y(v(y')) \\ &\leq (1-h)a + [h + (1-h)a]\delta b(\tilde{y}) + (1-h)(1-a)\delta b(y') \\ &= b(y) = B_Y(v), \end{aligned}$$

and $\tilde{z}, z'$ witness $z \in W_a(co[\text{graph}(B_Y)])$. In particular, $TB_Y(v) = TB_Y(v(z)) \leq b(z) \leq B_Y(v)$

Next, consider any $v \in \text{dom}(B_Y)$. There is some probability measure μ on the extreme points of Y such that $(v, B_Y(v)) = \int_Y y \, \mathrm{d}\mu(y)$. By minimality of $B_Y(v)$, it must be that $y \in \text{graph}(B_Y)$ a.s.-$\mu(y)$. So letting μ_1 be the marginal of μ on the first coordinate, $(v, B_Y(v)) = \int_{v(Y)} (u, B_Y(u)) \, \mathrm{d}\mu_1(u)$, so that

$$B_Y(v) = \int_{v(Y)} B_Y \, \mathrm{d}\mu_1 \geq \int_{v(Y)} TB_Y \, \mathrm{d}\mu_1 \geq TB_Y(v),$$

where the last inequality follows from Jensen's theorem.

This completes the proof.

□

The Bellman equation follows immediately.

Corollary 3. *$B := B_{\mathcal{E}}$ solves the Bellman equation $B = cvx\left[\min\left\{T_f B, \ \ T_0 B, \ \ T_1 B\right\}\right]$.*

Aligned Optimum In line with the main text, we now proceed to characterize the payoffs attainable in equilibria with no bad projects.

Notation. *Define the discrete time[36] marginal value of search, $\omega := \dfrac{h}{1-\delta}(\bar{\theta} - \underline{\theta})$.*

Before proceeding, we record the discrete time expression of Assumption 2. It is worth highlighting that, as $\Delta \to 0$, the continuous time version implies the discrete time one.

Assumption 2.

$\delta\omega \geq \underline{\theta}$ or, equivalently, $\omega \geq \theta_E$.

[36] Notice that, given parametrization $(\delta, h) = (1-\Delta, \eta\Delta)$, this coincides exactly with the definition from the main text.

Similarly to the main text, we assume Assumption 2 holds throughout the following analysis, unless otherwise stated. We now state our aligned equilibrium result.

Theorem 1 (Aligned Optimum)**.**

1. *There exist productive aligned equilibria if and only if Assumption 2 holds.*

2. *Every aligned equilibrium generates revenue* $\leq \omega$*.*

3. *Given Assumption 2,* $(\omega, 0) \in \mathcal{E}^*$.

Proof. We proceed in reverse. For (3), assumption 2 holds and define $B(v) = 0$ for $v \in [0, \omega]$ and $B(v) = \infty$ otherwise. Notice that, $TB(\omega) \leq TB_0(\omega) = \delta h B(\frac{\omega - exth}{\delta}) + \delta(1 - h)B(\frac{\omega - \Delta\omega}{\delta}) = 0$, where the first equality holds by Assumption 2. Because TB is convex, B is self-generating, and thus $B \geq B_{\mathcal{E}^*}$. (3) follows.

We now proceed to verify (2), i.e. that $\hat{v} > \omega$ implies that $(0, \hat{v}) \notin \mathcal{E}^*$.
Suppose $v > \omega$ has $B_{\mathcal{E}^*}(v) = 0$. Then $B_{\mathcal{E}^*}|_{[0,v]} = 0$, and

$$0 = B_{\mathcal{E}^*}(v) = TB_{\mathcal{E}^*}(v) = \min\{T_f B_{\mathcal{E}^*}(v), T_0 B_{\mathcal{E}^*}(v), T_1 B_{\mathcal{E}^*}(v)\}.$$

Notice that $TB_{\mathcal{E}^*}(v) \neq T_1 B_{\mathcal{E}^*}(v)$ as the latter is > 0. If $TB_{\mathcal{E}^*}(v) = T_f B_{\mathcal{E}^*}(v)$, then since $B_{\mathcal{E}^*}$ is increasing

$$B_{\mathcal{E}^*}\left(v + \frac{1-\delta}{\delta}(v - \omega)\right) \leq B_{\mathcal{E}^*}\left(v + \frac{1-\delta}{\delta}v\right) = \delta^{-1} T_f B_{\mathcal{E}^*}(v) = 0.$$

Finally, if $TB_{\mathcal{E}^*}(v) = T_0 B_{\mathcal{E}^*}(v)$, then[37]

$$0 = B_{\mathcal{E}^*}\left(\frac{v - (1-\delta)\omega}{\delta}\right) = \delta B_{\mathcal{E}^*}\left(v + \frac{1-\delta}{\delta}(v - \omega)\right).$$

So either way, $B_{\mathcal{E}^*}\left(v + \frac{1-\delta}{\delta}(v - \omega)\right) = 0$ too.
Now, if $\hat{v} > \omega$ with $(0, \hat{v}) \in \mathcal{E}^*$, then applying the above inductively yields a sequence[38] $v_n \to \infty$ on which $B_{\mathcal{E}^*}$ takes value zero. This would contradict the compactness of $B_{\mathcal{E}^*}$'s proper domain.

By (3) we know that productive aligned equilibria exist if Assumption 2 holds. To see the necessity of the assumption, suppose for a contradiction that it doesn't hold and yet

[37] Since a weighted average of two nonnegative numbers can be zero only if both numbers are zero.

[38] Let $v_0 = \hat{v}$ and $v_{n+1} = v_n + \frac{1-\delta}{\delta}(v_n - \omega) \geq \hat{v} + n(\hat{v} - \omega)$.

some productive aligned equilibrium exists. Let v be the agent's continuation value at some on-path history at which the principal delegates. By (2), we know $v \leq \omega$. As Assumption 2 fails, we then know $\delta v < \underline{\theta}$, contradicting agent IC. $\square$

Optimality Now, focus on the frontier of the whole equilibrium set. Before proceeding to the full characterization, we establish a single crossing result: indiscriminate project adoption is initially used only for the highest agent values.

Lemma 7. *Fix $B \in \mathcal{B}$, and suppose $B^{-1}(0) = [0, \omega]$:*

1. *If $v > \omega$, then $T_0 B(v) < T_f B(v)$ (unless both are ∞).*
2. *There is a cutoff $\underline{v} \geq \omega$ such that*

$$\begin{cases} T_0 B(v) \leq T_1 B(v) & \text{if } v \in [\omega, \underline{v}); \\ T_0 B(v) \geq T_1 B(v) & \text{if } v > \underline{v}. \end{cases}$$

Proof. $B(\omega) = 0$, and B is convex. Therefore, B is strictly increasing above ω on its domain, so that[39] $\check{T}_f B(v) > \check{T}_0 B(v)$, confirming the first point.

Given v,
$$\frac{v - [\theta_E - \underline{\theta}]}{\delta} - \frac{v - \theta_E}{\delta} = \delta^{-1}\underline{\theta}$$
is a nonnegative constant.

Since B is convex, it must be that the continuous function

$$v \mapsto \check{T}_0 B(v) - \check{T}_1 B(v) = \delta(1-h)\left[B\left(\frac{v - [\theta_E - \underline{\theta}]}{\delta}\right) - B\left(\frac{v - \theta_E}{\delta}\right)\right] - (1-h)$$

is increasing on its proper domain. The second point follows.[40] $\square$

Our characterization of the equilibrium frontier B can now be stated.

Theorem 2 (Equilibrium Frontier)**.**
Let $B := B_{\mathcal{E}^*}$ and $\bar{v} := \max dom(B)$.

1. *$\bar{v} \geq \omega$, and $B(v) = 0$ for $v \in [0, \omega]$.*

[39] The relationship is as shown if $\check{T}_0 B(v) < \infty$. Otherwise, $T_f B(v) = T_0 B(v) = \infty$.
[40] Because wherever $\check{T}_i B(v) \geq \check{T}_j B(v)$, we have $T_i B(v) \geq T_j B(v)$ as well.

2. *If $\bar{v} > \omega$, then*

$$\begin{aligned} B(v) &= T_0 B(v) \text{ for } v \in \left[\omega,\ \delta\bar{v} + (\theta_E - \underline{\theta})\right]; \\ B(v) &\text{ is affine in } v \text{ for } v \in \left[\delta\bar{v} + (\theta_E - \underline{\theta}),\ \bar{v}\right]; \\ B(\bar{v}) &= T_1 B(\bar{v}). \end{aligned}$$

3. *If $\bar{v} > \omega$, then $\pi(\bar{v}, B(\bar{v})) = 0$.*

Proof. The first point follows directly from Theorem 1. Now suppose $\bar{v} > \omega$.
Let $\underline{v} := \delta\bar{v} + (\theta_E - \underline{\theta})$. Any $v > \underline{v}$ has $\frac{v-[\theta_E-\underline{\theta}]}{\delta} > \frac{\underline{v}-[\theta_E-\underline{\theta}]}{\delta} = \bar{v}$, so that $B\left(\frac{v-[\theta_E-\underline{\theta}]}{\delta}\right) = \infty$, and therefore (appealing to Proposition 4) $T_0B(v) = \infty$. Therefore, the cutoff defined in Lemma 7 is $\leq \underline{v}$.

Since T_0B, T_1B are both convex, there exist some $v_0, v_1 \in [\omega, \bar{v}]$ such that $v_0 \leq v_1, \underline{v}$, and:

$$\begin{aligned} B(v) &= 0 \text{ for } v \in [0, \omega]; \\ B(v) &= T_0B(v) \text{ for } v \in [\omega, v_0]; \\ B(v) &\text{ is affine in } v \text{ for } v \in [v_0, v_1]; \\ B(v) &= T_1B(v) \text{ for } v \in [v_1, \bar{v}]. \end{aligned}$$

Let $m > 0$ denote the left-sided derivative of B at v_1 (which is simply the slope of B on (v_0, v_1) if $v_0 \neq v_1$).

Let $[v_0, v_1]$ be maximal (w.r.t. set inclusion) such that the above decomposition is still correct.

Notice then that $v_1 = \bar{v}$. Indeed:

- If $\frac{v_1-\theta_E}{\delta} \geq v_0$, then (appealing to Proposition 4) the right-side derivative $T_1B' = m$ in some neighborhood of v_1 in $[v_1, \bar{v}]$. By convexity of B, this would imply $B(v) = T_1B(v) = B(v_1) + m(v - v_1)$ in said neighborhood. Then, by maximality of v_1, it must be that $v_1 = \bar{v}$.

- If $\frac{v_1-\theta_E}{\delta} < v_0$, then minimality of v_0 implies (again using Proposition 4) that the right-sided derivative $B'(v_1) = TB'(v_1) < m$ if $v_1 < \bar{v}$. As B is convex, this can't happen. Therefore, $v_1 = \bar{v}$.

Finally, we need to show that $v_0 = \underline{v}$. Now, by minimality of v_0, it must be that for any $v \in [0, v_0)$, the right-side derivative $B'(v) < m$. If $v_0 < \underline{v}$ (so that $\frac{v_0-[\theta_E-\underline{\theta}]}{\delta} < \underline{v}$), then

Proposition 4 gives us

$$\begin{aligned} m &= B'(v_0) \\ &\leq T_0 B'(v_0) \\ &= hB'\left(\frac{v_0 - \theta_E}{\delta}\right) + (1-h)B'\left(\frac{v_0 - [\theta_E - \underline{\theta}]}{\delta}\right) \\ &= hB'\left(\frac{v_0 - \theta_E}{\delta}\right) + (1-h)m \\ &< m, \end{aligned}$$

a contradiction. Therefore $v_0 = \underline{v}$, and the second point of the theorem follows.

For the last point, assume $\bar{v} > \omega$ and yields strictly positive profits. Then, for sufficiently small $\gamma > 0$, the function $B^\gamma \in \mathcal{B}$ given by $B^\gamma(v) = \begin{cases} B(v) & \text{if } v \in [0, \bar{v}] \\ T_1 B(v) & \text{if } v \in [\bar{v}, \bar{v} + \gamma] \end{cases}$ is self-generating, contradicting the fact that $\mathcal{E}$ is the largest self-generating set. □

1.5 The Efficient Frontier: Continuous Time

In this section we proceed to take the limit of the discrete time equilibrium value set frontier. We consider the limit with $(\delta, h) = (\delta_\Delta, h_\Delta) := (1-\Delta, \eta\Delta)$ (and so $\theta_E = \underline{\theta} + \omega\Delta$) as the period length $\Delta \to 0$.

Let B_Δ be the efficient frontier of the equilibrium value set. Given Assumption 2, for Δ is sufficiently small (so that the discrete time version of Assumption 2 holds as well), Theorem 2 applies:

$$\begin{aligned} B_\Delta(v) &= 0 \text{ for } v \in [0, \ \omega]\,; \\ B_\Delta(v) &= (1-\Delta)\left[\eta\Delta B_\Delta\left(\frac{\bar{v}_\Delta - \underline{\theta} - \Delta\omega}{1-\Delta}\right) + (1-\eta\Delta)B_\Delta\left(\frac{\bar{v}_\Delta - \Delta\omega}{1-\Delta}\right)\right] \text{ for } v \in [\omega, \ (1-\Delta)\bar{v}_\Delta + \Delta\omega]\,; \\ B_\Delta(v) &\text{ is affine in } v \text{ for } v \in [(1-\Delta)\bar{v}_\Delta + \Delta\omega, \ \bar{v}_\Delta]\,; \\ B_\Delta(\bar{v}) &= (1-\eta\Delta) + (1-\Delta)B_\Delta\left(\frac{\bar{v}_\Delta - \underline{\theta} - \Delta\omega}{1-\Delta}\right). \end{aligned}$$

Define D_0^Δ and F_0^Δ on $\mathcal{B}$ via:

$$\begin{aligned}
D_0^\Delta B &:= \frac{1}{\Delta}\left[B - T_{\Delta,f}^{-1} T_{\Delta,0} B\right], \text{ and} \\
F_0^\Delta B &:= \frac{1}{\Delta}\left[B - T_{\Delta,f}^{-1} B\right].
\end{aligned}$$

Notice that

$$\begin{aligned}
\Delta F_0^\Delta B(v) &= B(v) - \eta\Delta B(v-\theta_E) - (1-\eta\Delta)B\left(v - [\theta_E - \underline{\theta}]\right) \\
&= B(v) - \eta\Delta B(v-\underline{\theta}-\Delta\omega) - (1-\eta\Delta)B(v-\Delta\omega) \\
&= \eta\Delta\left[B(v) - B(v-\underline{\theta}-\Delta\omega)\right] + (1-\eta\Delta)\left[B(v) - B(v-\Delta\omega)\right] \\
\implies F_0^\Delta B(v) &= \eta\left[B(v) - B(v-\underline{\theta}-\Delta\omega)\right] + (1-\eta\Delta)\frac{B(v) - B(v-\Delta\omega)}{\Delta}, \text{ and} \\
D_0^\Delta B(v) &= \frac{B(v) - T_{\Delta,f}^{-1} B(v)}{\Delta} \\
&= \frac{B(v) - \frac{1}{1-\Delta} B((1-\Delta)v)}{\Delta} \\
&= \frac{1}{1-\Delta}\left[\frac{B(v) - B(v-\Delta v)}{\Delta} - B(v)\right], \\
\implies F_0^\Delta B(v) - D_0^\Delta B(v) &= \left(\frac{1}{1-\Delta} + \eta\right)B(v) - \left[\eta B(v-\underline{\theta}-\Delta\omega) + \frac{B(v) - B(v-\Delta\omega)}{(1-\eta\Delta)^{-1}\Delta} - \frac{B(v) - B(v-\Delta v)}{(1-\Delta)\Delta}\right].
\end{aligned}$$

For fixed $\Delta > 0$, letting $B = B_\Delta$, and taking $v \in [\omega, \bar{v}_\Delta]$, the LHS is zero. For fixed B, in the limit as $\Delta \to 0$, the RHS converges to

$$(1+\eta)B(v) - [\eta B(v-\underline{\theta}) + (v-\omega)B'(v)].$$

If $\bar{v} = \lim_{\Delta\to 0} \bar{v}_\Delta$ and $B = \lim_{\Delta\to 0} B_\Delta$ exist, then

$$\begin{aligned}
B(v) &= \begin{cases} 0 & \text{if } v \in [0,\omega] \\ \dfrac{\eta B(v-\underline{\theta}) + (v-\omega)B'(v)}{1+\eta} & \text{if } v \in (\omega, \bar{v}) \\ 1 + B(v-\underline{\theta}) & \text{if } v = \bar{v} > \omega, \end{cases} \\
\bar{v} &= \max \operatorname{dom}(B) \\
\pi(\bar{v}, B(\bar{v})) &= 0.
\end{aligned}$$

Summarizing, there is some $\bar{v} \geq \omega$ such that:

- The highest revenue attainable in equilibrium is $\bar{v}$.
- An optimal way to provide revenues $v \in [0, \omega]$ in equilibrium is with shutdown of duration $\log \frac{\omega}{v}$ followed by the Aligned Optimal Budget.
- If $\bar{v} > \omega$, then
  - The optimal way to provide revenue $v \in (\omega, \bar{v})$ in equilibrium is delegation with picky project adoption, jumping to continuation revenue $v - \underline{\theta}$ following a project, and revenue drifting according to $\dot{v} = v - \omega$ conditional on no projects.
  - The only way to provide revenue $\bar{v}$ in equilibrium is with an immediate bad project followed by providing revenue $\bar{v} - \underline{\theta}$.
  - The optimal equilibrium profit from providing $\bar{v}$ is zero.

2 APPENDIX: Expiring Budget

Consider the cyclic contract with fiscal year of length $T \in (0, \infty)$. Rather than working with T as our parameter, it is convenient to work with the transformed variable

$$z = \sum_{k=0}^{\infty} e^{-kT} = \frac{1}{1 - e^{-T}} > 1, \text{ so that } e^{-T} = \frac{z-1}{z}.$$

The expected discounted number of bad/good projects are

$$\begin{aligned} b(z) &= ze^{-T}e^{-\eta T} = (z-1)\left(\frac{z-1}{z}\right)^{\eta}, \text{ and} \\ g(z) &= z\int_0^T e^{-t}\eta e^{-\eta t}\,\mathrm{d}t = \frac{z\eta}{1+\eta}[1 - e^{(1+\eta)T}] \\ &= \frac{\eta}{1+\eta}(z-b). \end{aligned}$$

Which $z > 1$ maximizes $\mathcal{P}$'s profit? Letting $\rho := \frac{\bar{\theta}-c}{c-\underline{\theta}} > 0$, profit is proportional to

$$\rho g - b = \frac{\rho\eta}{1+\eta}(z-b) - b \propto \frac{\rho\eta}{\rho\eta + 1 + \eta}z - b.$$

It is straightforward to show that profit is uniquely[41] maximized at $z^* = z(\eta)$, the unique value $z^* > 1$ with[42]

$$b'(z^*) = \frac{\rho\eta}{\rho\eta + 1 + \eta}.$$

How high is the profit, as a fraction of first-best profit? Some algebra shows that

$$\begin{aligned} \frac{\text{profit from optimal Expiring Budget contract}}{\text{first-best profit}} &= \frac{\rho g - b}{\rho\eta - 0} \\ &= \frac{1}{1+\ell}, \end{aligned}$$

where $\ell(\eta) := \frac{\eta}{z(\eta)}$. Then,

$$\left(1 + \frac{-\ell}{\eta}\right)^{\eta}(1+\ell) = b'(\eta) = \frac{\rho\eta}{\rho\eta + 1 + \eta} \xrightarrow{\eta\longrightarrow\infty} \frac{\rho}{\rho+1} = \frac{\bar{\theta} - c}{\bar{\theta} - \underline{\theta}}.$$

Therefore, $\ell^* := \lim_{\eta\to\infty} \ell(\eta)$ exists and is the unique positive number satisfying $e^{-\ell^*}(1 + \ell^*) = \dfrac{\bar{\theta} - c}{\bar{\theta} - \underline{\theta}}$. In particular $\ell^* > 0$ (and so $\frac{1}{1+\ell^*} < 1$) so that, even as the players become arbitrarily patient, the Expiring Budget contract cannot approximate efficiency. For some parameter values, the best Expiring Budget contract outperforms the aligned optimum. Indeed, for any $\alpha, \beta \in (0,1)$, there are parameter values[43] $\bar{\theta} > c > \underline{\theta} > 0$ such that $\frac{\bar{\theta}-\underline{\theta}}{\bar{\theta}} = \alpha$ and $\frac{1}{1+\ell^*} = \beta$.

3 APPENDIX: Delayed Differential Equation

Taking a change of variables, from agent value v to account balance $x = \frac{v-\omega}{\underline{\theta}}$, the following system of equations describes the frontier of the equilibrium value set.

[41] Uniqueness is guaranteed by the convexity of b, together with a simple analysis of the limiting cases.

[42] The given z^* describes an equilibrium if and only if the principal has nonnegative continuation profit at the end of a project cycle. Given the below work comparing profit to first-best profit, this condition holds when η is sufficiently high.

[43] For instance, take $\bar{\theta} = 1$, $\underline{\theta} = 1 - \alpha$, and $c = 1 - \frac{\alpha}{\beta} e^{-\frac{1-\beta}{\beta}}$.

$$\begin{aligned}(1+\eta)b(x) &= \eta b(x-1) + xb'(x) \text{ for } x > 0\\ b(x) &= 0 \text{ for } x \le 0\end{aligned}$$

Theorem 4. *Consider the above system of equations. For any $\alpha \in \mathbb{R}$, there is a unique solution $b^{(\alpha)}$ to the above system with $b^{(\alpha)}(1) = \alpha$. Moreover $b^{(\alpha)} = \alpha b^{(1)}$.*
Letting $b = b^{(1)}$:

1. *$b(x) = x^{1+\eta}$ for $x \in [0, 1]$.*

2. *b is twice-differentiable on $(0, \infty)$ and globally C^1 on $\mathbb{R}$.*

3. *b is strictly convex on $(0, \infty)$ and globally convex on $\mathbb{R}$. In particular, b is unbounded.*

4. *b is strictly increasing and strictly log-concave on $(0, \infty)$.*

Proof. First consider the same equation on a smaller domain,

$$(1+\eta)b(x) = xb'(x) \text{ for } x \in (0, 1].$$

As is standard, the full family of solutions is $\{b^{(\alpha,1)}\}_{\alpha\in\mathbb{R}}$. where $b^{(\alpha,1)}(x) = \alpha x^{1+\eta}$ for $x \in (0, 1]$.

Now, given a particular partial solution $b : (-\infty, z] \to \mathbb{R}$ up to $z > 0$, there is a unique solution to the first-order linear differential equation $\hat{b} : [z, z+1] \to \mathbb{R}$ given by

$$\hat{b}'(x) = \frac{1+\eta}{x}\hat{b}(x) - \frac{\eta}{x}b(x-1).$$

Proceeding recursively, there is a unique solution to the given system of equations for each α. Moreover, since multiplying any solution by a constant yields another solution, uniqueness implies $b^{(\alpha)} = \alpha b^{(1)}$. Now let $b := b^{(1)}$.

We have shown that $b(x) = x^{1+\eta}$ for $x \in [0, 1]$, from which it follows readily that b is $C^{1+\lfloor\eta\rfloor}$ on $(-\infty, 1)$, Given $x > 0$, for small ϵ,

$$
\begin{aligned}
(x+\epsilon)\frac{b'(x+\epsilon)-b'(x)}{\epsilon} &= \frac{1}{\epsilon}(x+\epsilon)b'(x+\epsilon) - \frac{1}{\epsilon}xb'(x) - b'(x) \\
&= \frac{1}{\epsilon}\Big[(1+\eta)b(x+\epsilon) - \eta b(x+\epsilon-1)\Big] - \frac{1}{\epsilon}\Big[(1+\eta)b(x) - \eta b(x-1)\Big] - b'(x) \\
&= \eta\left[\frac{b(x+\epsilon)-b(x)}{\epsilon} - \frac{b(x-1+\epsilon)-b(x-1)}{\epsilon}\right] + \left[\frac{b(x+\epsilon)-b(x)}{\epsilon} - b'(x)\right] \\
&\xrightarrow{\epsilon\to 0} \eta[b'(x) - b'(x-1)] + 0.
\end{aligned}
$$

So b is twice differentiable at $x > 0$ with $b''(x) = \dfrac{\eta}{x}[b'(x) - b'(x-1)]$.

Let $\bar{x} := \sup\{x > 0 : \ b'|_{(0,x]} \text{ is strictly increasing}\}$. We know $\bar{x} \geq 1$, from our explicit solution of b up to 1. If $\bar{x}$ is finite, then $b'(\bar{x}) > b'(\bar{x}-1)$. But then $b''(\bar{x}) = \frac{\eta}{x}[b'(x) - b'(x-1)] > 0$, so that b' is strictly increasing in some neighborhood of $\bar{x}$, contradicting the maximality of $\bar{x}$. So $\bar{x} = \infty$, and our convexity result obtains. From that and $b'(0) = 0$, it is immediate that b is strictly increasing on $(0, \infty)$.

Lastly, let $f := \log b|_{(0,\infty)}$. Then $f(x) = (1+\eta)\log x$ for $x \in (0, 1]$, and for $x \in (1, \infty)$,

$$
\begin{aligned}
(1+\eta)e^{f(x)} &= \eta e^{f(x-1)} + xe^{f(x)}f'(x), \\
\implies (1+\eta) &= \eta e^{f(x-1)-f(x)} + xf'(x). \\
\implies 0 &= \eta e^{f(x-1)-f(x)}[f'(x-1) - f'(x)] + f'(x) + f''(x) \\
\implies -f''(x) &= \eta e^{f(x-1)-f(x)}[f'(x-1) - f'(x)] + f'(x) \\
&\geq \eta e^{f(x-1)-f(x)}[f'(x-1) - f'(x)], \text{since } f = \log b \text{ is increasing.}
\end{aligned}
$$

The same contagion argument will work again. If f has been strictly concave so far, then $f'(x) < f'(x-1)$, in which case $-f''(x) > 0$ and f will continue to be concave. Since we know $f|_{(0,1]}$ is strictly concave, it follows that f is globally such. □

The first point of the following proposition shows that the economically intuitive boundary condition of our DDE uniquely pins down the solution b for any given account cap. The second point shows that as the account cap increases, so does the number of bad projects (in expected discounted terms) anticipated at the cap.

Proposition 5. *For any $\bar{x} > 0$*

1. *There is a unique $\alpha = \alpha(\bar{x}) > 0$ such that $b^{(\alpha)}(\bar{x}) = 1 + b^{(\alpha)}(\bar{x}-1)$.*

2. *$b^{\alpha(\bar{x})}(\bar{x})$ is increasing in $\bar{x}$.*

Proof. The first part is immediate, with $\alpha = \frac{1}{b^{(1)}(\bar{x}) - b^{(1)}(\bar{x}-1)}$.

For the second part, notice that $\frac{b(\bar{x})}{b(\bar{x}-1)}$ is decreasing in $\bar{x}$ because b is log-concave. Then,

$$b^{\alpha(\bar{x})}(\bar{x}) = \alpha(\bar{x})b(\bar{x}) = \frac{b(\bar{x})}{b(\bar{x}) - b(\bar{x}-1)} = \frac{1}{1 - \frac{b(\bar{x}-1)}{b(\bar{x})}}$$

is increasing in $\bar{x}$. □

4 APPENDIX: Comparative Statics

In this section, we prove Proposition 3.

For any parameters $\eta, \bar{\theta}, \underline{\theta}, c$ satisfying Assumptions 1 and 2, and for any balance and bad projects x, b satisfying $x \geq b > 0$, define the associated profit

$$\begin{aligned}
\hat{\pi}(x, b|\eta, \bar{\theta}, \underline{\theta}, c) &:= \eta(\bar{\theta} - \underline{\theta}) + \underline{\theta}x - c\left[\frac{\eta(\bar{\theta} - \underline{\theta}) + \underline{\theta}x - \underline{\theta}b}{\bar{\theta}} + b\right] \\
&= \left(1 - \frac{c}{\bar{\theta}}\right)[\eta(\bar{\theta} - \underline{\theta}) + \underline{\theta}x] - c\left(1 - \frac{\underline{\theta}}{\bar{\theta}}\right)b \\
&= (\bar{\theta} - c)\eta + \left(1 - \frac{c}{\bar{\theta}}\right)\underline{\theta}(x - \eta) - c\left(1 - \frac{\underline{\theta}}{\bar{\theta}}\right)b.
\end{aligned}$$

For reference, we compute the following derivatives of profit:

$$\begin{aligned}
\frac{\partial\hat{\pi}}{\partial\bar{\theta}} &= \eta - c\underline{\theta}\frac{-1}{\bar{\theta}^2}[x - b - \eta] = \left(1 - \frac{c\underline{\theta}}{\bar{\theta}^2}\right)\eta + \frac{c\underline{\theta}}{\bar{\theta}^2}(x - b) > 0. \\
\frac{\partial\hat{\pi}}{\partial\underline{\theta}} &= \left(1 - \frac{c}{\bar{\theta}}\right)(x - \eta) + c\frac{1}{\bar{\theta}}b, \text{ which implies} \\
(\bar{\theta} - \underline{\theta})\frac{\partial\hat{\pi}}{\partial\underline{\theta}} + \hat{\pi} &= (\bar{\theta} - c)\eta + \left(1 - \frac{c}{\bar{\theta}}\right)\bar{\theta}(x - \eta) - 0b = (\bar{\theta} - c)x > 0. \\
\frac{\partial\hat{\pi}}{\partial c} &= -\frac{1}{\bar{\theta}}\left[\eta(\bar{\theta} - \underline{\theta}) + \underline{\theta}x + (\bar{\theta} - \underline{\theta})b\right] < 0.
\end{aligned}$$

Fix parameters $(\bar{\theta}, \underline{\theta}, c)$, and let $\bar{x}^*$ be as delivered in Theorem 2. We first show that slightly raising either of $\bar{\theta}, \underline{\theta}$ or slightly lowering c weakly raises the cap, strictly if $\bar{x}^* > 0$.

- If $\bar{x}^* = 0$, there is nothing to show, so assume $\bar{x}^* > 0$ henceforth. Notice that the expected discounted number of bad projects when at the cap depends only on η and the size of the cap. By Theorem 2, $\hat{\pi}(\bar{x}^*, b|\eta, \bar{\theta}, \underline{\theta}, c) = 0$.

- Consider slightly raising $\bar{\theta}$ or $\underline{\theta}$, or slightly lowering c. By the above derivative computations, the profit of the DCB contract with cap $\bar{x}^*$ is strictly positive.

- For any of the above considered changes, the DCB contract with cap $\bar{x}^*$ has strictly positive profits. Appealing to continuity, a slightly higher cap still yields positive profits under the new parameters, and is therefore consistent with equilibrium by Proposition 2. Then, appealing to Theorem 2 again, the cap associated with the new parameters is strictly higher than $\bar{x}^*$.

Now we consider comparative statics in the profit-maximizing initial account balance. We will fix parameters $(\bar{\theta}, \underline{\theta}, c)$, and consider raising either of $\bar{\theta}, \underline{\theta}$ or lowering c.

- Again, if $\bar{x}^* = 0$ at the original parameters, there's nothing to check, so assume $\bar{x}^* > 0$.

- Let $\check{b}$ be some solution to the DDE in Section 3, so that the expected discounted number of bad projects at a given cap $\bar{x}$ and balance x is $b(x|\bar{x}) = \dfrac{\check{b}(x)}{\check{b}(\bar{x}) - \check{b}(\bar{x}-1)}$. Because $\check{b}$ is strictly increasing and strictly convex (by work in Section 3), we know that $\frac{\partial}{\partial x} b(x|\bar{x})$ is strictly decreasing in $\bar{x}$. Therefore, by our comparative statics result for the cap, the parameter change results in a global strict decrease of $\frac{\partial}{\partial x} b(x|\bar{x}^*)$.

- By the form of $\hat{\pi}$ and by convexity of $b(\cdot|\bar{x})$, the unique optimal initial balance is the balance at which $\frac{\partial}{\partial x} b(x|\bar{x})$ is equal to

$$\xi = \frac{\underline{\theta}\left(1 - \frac{c}{\bar{\theta}}\right)}{c\left(1 - \frac{\underline{\theta}}{\bar{\theta}}\right)} = \frac{\underline{\theta}(\bar{\theta} - c)}{c(\bar{\theta} - \underline{\theta})},$$

which increases with the parameter change.

- As ξ increases and $\frac{\partial}{\partial x} b(x|\bar{x}^*)$ decreases (at each x) with the parameter change, the optimal balance x^* must increase (given convexity) to satisfy the first-order condition $\left.\frac{\partial}{\partial x}\right|_{x=x^*} b(x|\bar{x}) = \xi$.